

Accelerating Structure Discovery for Functional Materials



Jordan Antony Dorrell

Supervisor: Prof. A. J. Morris

Metallurgy and Materials

University of Birmingham

This dissertation is submitted for the degree of

PhD

September 2023

UNIVERSITY OF
BIRMINGHAM

University of Birmingham Research Archive

e-theses repository

This unpublished thesis/dissertation is copyright of the author and/or third parties. The intellectual property rights of the author or third parties in respect of this work are as defined by The Copyright Designs and Patents Act 1988 or as modified by any successor legislation.

Any use made of information contained in this thesis/dissertation must be in accordance with that legislation and must be properly acknowledged. Further distribution or reproduction in any format is prohibited without the permission of the copyright holder.

Abstract

Discovering new functional materials and furthering understanding of known functional materials benefits from computational structure searching approaches that can identify the ground state structure for a particular system. To this end, using both density-functional theory (DFT) and machine learning (ML) models, I have searched for the structures of new materials with potential applications as electrodes for alkali-ion batteries and as single-photon emitters (SPEs).

To search for new Li-Ni-S ternary electrodes with a high gravimetric charge capacity, I have taken a random structure searching approach. Initial searches were performed with an *ab initio* model, but the process was accelerated by three orders of magnitude by switching to a machine learning model. Over 150,000 structures were generated and while no ternary compounds were identified on the convex hull, new polymorphs of NiS and NiS₂ were discovered that are lower in formation energy than the structures reported experimentally and on the Materials Project.

Ab initio random structure searching (AIRSS) was also applied to search for point defects responsible for single photon emission in monolayer hBN. Previous studies of hBN as an SPE have focused on only a limited range of intrinsic defects. In this work I have generated over 5000 intrinsic and extrinsic defects and calculated their defect formation energies. Numerous new low energy defects were discovered that have not previously been considered as SPE sources. The optical absorption and emission spectra of these defects can now be modelled

so that they can be compared to experiment in hopes of identifying defects that give rise to desirable optical phenomena.

K-intercalated graphite (KGIC) is a potential anode material for K-ion batteries and its electrochemical synthesis can be observed unobtrusively through Raman spectroscopy. To better our understanding of the relationship between the structure and the Raman spectra, I have modelled Raman spectra for stages 1-5 KGIC within the conventional Rüdorff-Hofmann (RH) model and a range of candidate cells within the Daumas-Hérolde (DH) model, which were generated with a bespoke machine-learning driven algorithm. DH-type cells were energetically accessible at room temperature but the modelled Raman spectra showed poor correlation with experiment. I conclude that the RH model is more consistent with experimentally observed KGICs, but the material may contain DH-type defects in small quantities.

In addition, two Na-niobates with potential electrode applications, layered Na_3NbO_4 and $\text{NaNb}_7\text{O}_{18}$, were recently synthesised. To verify the structures of these compounds, I have modelled their ^{24}Na NMR spectra. In the case of $\text{NaNb}_7\text{O}_{18}$, the positions of the Na^+ ions were unknown and so I have modelled the spectra of a range of candidate cells. The total energy of these candidate cells differed by less than 3.3 meV / atom, suggesting that the Na^+ ions are disorderly distributed between several sites throughout the structure. The modelled ^{24}Na NMR spectrum of layered Na_3NbO_4 shows exceptional correlation with the experimental spectrum, suggesting a strong match between the candidate structure and the synthesised material. The modelled ^{24}Na NMR spectrum of $\text{NaNb}_7\text{O}_{18}$ contained three peaks which correlate well with peaks in the experimental spectrum. The modelled spectra did not reproduce the fourth peak, which has been attributed to a NaNb_3O_8 impurity. The model cells of $\text{NaNb}_7\text{O}_{18}$ therefore also seem to be consistent with the experimental sample.

Acknowledgements

I would like to thank for my supervisor, Andrew Morris, for his support and guidance during this research. For their useful discussions, I also thank the members of the Morris Group and the Arhangelskis Group, namely, Hrishit Banerjee, James Darby, Matthew Evans, Wei Huang, Annalena Genreith-Schriever, Angela Harper, Andrea Iliceto, Can Koçer, Mario Ongkiko, Chris Owen, Wataru Sekine, Weronika Wiesiołek, Mihails Arhangelskis, and Yizhi Xu. This research was made possible by the contributions made by my collaborators, Mauro Pasta, Isaac Capone, Shobhan Dhir, Bartomeu Monserrat, Siyu Chen, and Kent Griffith. I thank the University of Birmingham's BEAR team who maintain and support the local cluster computer, BlueBEAR, which has reliably afforded me compute time for the full duration of my research. Computing has also generously been provided by Baskerville at the University of Birmingham's, Athena and Sulis through HPC Midlands+'s, CSD3 and Cluster2 at the University of Cambridge, Bede at the University of Durham, Prometheus through PL-Grid, and SDCC at Brookhaven National Laboratory. Lastly, I gratefully acknowledge funding from the EPSRC.

Table of contents

List of figures	xi
List of tables	xv
1 Introduction	1
2 Methods	11
2.1 Structure Prediction	11
2.1.1 Potential Energy Surfaces	12
2.1.2 Geometry Optimisations	16
2.1.3 <i>Ab Initio</i> Random Structure Searching	23
2.1.4 Other Structure Searching Algorithms	25
2.2 Kohn-Sham DFT	29
2.2.1 Exchange-Correlation Functionals	33
2.2.2 Semi-Empirical Dispersion Corrections	34
2.2.3 Pseudopotentials	35
2.2.4 Bloch Theorem and Plane waves	37
2.2.5 DFT in Practice	40
2.3 Artificial Neural Networks	42
2.3.1 Early Architecture	42
2.3.2 Training	43

2.3.3	Atomistic ANNs	46
2.3.4	ANNs in this Project	48
2.4	Observables	51
2.4.1	Phonons	52
2.4.2	Vibrational Spectroscopy	55
2.4.3	NMR	56
3	LiNiS Electrodes	61
3.1	Introduction	61
3.2	Ternary Phase Diagram	63
3.3	Structure Searching	71
3.3.1	Element Swapping	71
3.3.2	AIRSS and MLRSS	73
3.3.3	Element Swapping with AIRSS	78
3.3.4	Higher Accuracy Geometry Optimisations	79
3.3.5	Ni-S Binary Convex Hull	81
3.3.6	Li-Ni-S Ternary Convex Hull	84
3.3.7	NiS ₂ -Li Tie Line	86
3.4	Conclusion	89
4	Monolayer hBN Defects	91
4.1	Introduction	91
4.2	Defect Searching	94
4.3	Defect Formation Energies	97
4.4	Low Energy Structures	101
4.5	Conclusion	106

5	Potassium Graphite Intercalation Compounds	109
5.1	Introduction	109
5.2	Daumas-Hérolde Defects	113
5.2.1	Introduction	113
5.2.2	Structure Searching Method	115
5.2.3	Structure Analysis	121
5.3	Electronic Structure	124
5.4	Phonon and Raman Spectra	126
5.5	Conclusion	132
6	Niobate Cathodes	137
6.1	Introduction	137
6.2	X_3NbO_4	140
6.3	$NaNb_7O_{18}$	144
6.3.1	Structure Searching Method	144
6.3.2	Structure Analysis	148
6.3.3	NMR	151
6.4	Conclusion	155
7	Conclusion	159
8	Further Work	167
9	Supplemental Data	171
	References	177

List of figures

1.1	Energy Consumption in England and Wales	3
2.1	2-Atom Cluster Lennard-Jones PES	14
2.2	3-Atom Cluster Lennard-Jones PES	15
2.3	1-Atom Crystal Lennard-Jones PES	17
2.4	AIRSS and Simulated Annealing Algorithms	26
2.5	Molecular Dynamics Example	26
2.6	Nested Sampling Algorithm	27
2.7	GA and PSO Algorithms	30
2.8	Two-Layer Perceptron	44
2.9	Convolution Example	50
3.1	LiMO ₂ Unit Cells	62
3.2	Example Binary Convex Hull	67
3.3	MLRSS Flow Chart	77
3.4	LiNiS Learning Curve	80
3.5	Difference in Formation Energy after Higher Accuracy Geometry Optimisations	80
3.6	NiS Binary Convex Hull	82
3.7	NiS Binary Minima	82
3.8	Ni ₃ S ₂ Polymorphs within 26 meV of Hull	83

3.9	NiS Polymorphs within 26 meV of Hull	85
3.10	LiNiS Ternary Compound Closest to Convex Hull	85
3.11	LiNiS Ternary Convex Hull	86
3.12	NiS ₂ Conversion Electrode Voltage Curve	88
3.13	Li-NiS ₂ Binary Convex Hull Distance	88
4.1	hBN AIRSS Cell	95
4.2	hBN Defects	98
4.3	hBN Defect Ternary AIRSS Overview	99
4.4	hBN Convergence	101
4.5	Difference in Formation Energy Between Defects in hBN Cells of Different Sizes	102
4.6	hBN Defect Formation Energies	105
4.7	hBN Axial Defects	106
5.1	Stage 1-5 RH Cells	112
5.2	RH and DH Unit cell Schematics	113
5.3	Stage 4 DH Cell Geometry Optimisation	116
5.4	Base Cells for DH Construction	118
5.5	KGIC ANN Force Convergence	120
5.6	KGIC ANN Force Validation Error	121
5.7	Stage 2-5 DH Cells	122
5.8	Stage 4 DH Cell - Large	123
5.9	Electronic DOS for KGIC RH Cells	125
5.10	Electronic DOS for KGIC DH Cells	126
5.11	RH KGIC Phonon Spectra	128
5.12	RH KGIC Raman Spectra	129
5.13	KGIC Experimental Raman Spectra 532 nm	130

5.14	KGIC Raman Spectra G-Peak	132
6.1	LiCoO ₂ - Li ₃ NbO ₄ Comparison	138
6.2	NaNb ₇ O ₁₈ Unit Cell	141
6.3	Li ₃ NbO ₄ Cells	141
6.4	Na ₃ NbO ₄ Cells	142
6.5	Li ₃ NbO ₄ Chemical Shieldings	142
6.6	Na ₃ NbO ₄ Chemical Shieldings	143
6.7	Layered Na ₃ NbO ₄ Experimental NMR Spectra	145
6.8	NaNbO Convergence	147
6.9	NaNb ₇ O ₁₈ Chemical Shieldings	149
6.10	NaNb ₇ O ₁₈ Candidate Unit Cells	150
6.11	NaNb ₇ O ₁₈ with 8 Formula Units	151
6.12	NaNb ₇ O ₁₈ Supercell Comparison	154
6.13	NaNb ₇ O ₁₈ PDF Comparison	154
6.14	Layered NaNb ₇ O ₁₈ Experimental NMR Spectra	157
9.1	Normalised Electronic DOS for KGIC RH Cells	174
9.2	Effect of choice of chemical potential on defect energies	175
9.3	hBN Axial Defects	176
9.4	L-Na ₃ NbO ₄ Environments	176

List of tables

5.1	LiNiS Initial GO Parameters	135
5.2	Mulliken Analysis of KGIC RH Cells	136
9.1	LiNiS Initial GO Parameters	171
9.2	LiNiS ANN hyperparameters	171
9.3	LiNiS Higher Quality GO Parameters	172
9.4	KGIC DFT Parameters	172
9.5	KGIC ANN hyperparameters	172
9.6	hBN geometry optimisation parameters	173
9.7	NbO geometry optimisation and NMR parameters	173

Chapter 1

Introduction

Society has demanded progressively more energy since the first industrial revolution in the 1700's, as shown in Figure 1.1. In the 18th century most energy in England and Wales was generated through the combustion of coal or wood and transferred through mechanical power [1]. The way that we both produce and consume energy has subsequently changed dramatically, firstly via the transition from mechanical to electrical power transfer mechanisms in the second industrial revolution, but also more recently in the increased use of handheld/wireless devices/vehicles and the transition to renewable energy sources. Some renewable energy sources, such as wind and solar do not provide consistent power, and may therefore not always meet demands. One possible solution is to invest in enough energy-generation infrastructure that the minimum power supply exceeds the maximum power demand but this is highly expensive and inefficient. Instead, it is preferable to store energy when supply exceeds demand, and then pull from that energy store when demand later exceeds supply. A range of energy storage solutions exist but not the practicality of each solution depends on the circumstances. Pumped-storage hydroelectricity is highly efficient and eco-friendly but it requires specific geological features and it usually requires the construction of large and expensive infrastructure. Fossil fuels, however, have been widely implemented as a means of generating energy *in situ* in vehicles for over two centuries but

are now falling out of favour due to their negative impact on the environment. Fossil fuels are also impractical for handheld devices due to the generation of heat, sound, and gaseous byproducts. H fuel cells are an emerging technology that has the advantage over fossil fuels that the only byproduct is water, which is more environmentally friendly and can be recycled back into H to “recharge” the fuel cell. It is also possible to exploit energy stored in nuclei via fission reactors and radioisotope generators. Like fossil fuels, the fuel must be harvested from natural sources and therefore is considered non-renewable. Nuclear energy sources produce consistent power and have a very long lifespan but they are expensive and the radiation poses a threat to human health. Mobile fission reactors are usually therefore only found on military vessels and radioisotope generators are most famously used on satellites. Electrochemical batteries exploit the difference in redox potential between two electrodes to generate an electrical current. They are highly practical in that energy is output from the device, and in the case of rechargeable batteries, input to the device in the form of an electrical current: no intermediate mechanism is needed to transform energy from one form to another. They are also environmentally friendly (though the impact from manufacturing/disposal is usually significant, depending on the battery materials), and highly scalable. Electrochemical batteries are therefore the most popular choice for handheld devices, they are frequently used for large-scale energy storage by nation-states, and they are also becoming widely implemented for terrestrial vehicles such as cars and light delivery vehicles, with the battery demand for EVs increasing by an order of magnitude from 43.8 GWh/year in 2016 to 550.5 GWh/year in 2022 [2].

Li-ion batteries (LIBs) were first commercialised in 1991 and have since been widely implemented around the globe [3] in a variety of common commercial and industrial devices such as mobile phones and electric vehicles. Goodenough’s original LIBs used a graphite anode, and a LiCoO_2 cathode which were separated by an electrolyte: a medium that allowed for the transfer of ions but not lone electrons. Each of these electrodes were then mounted

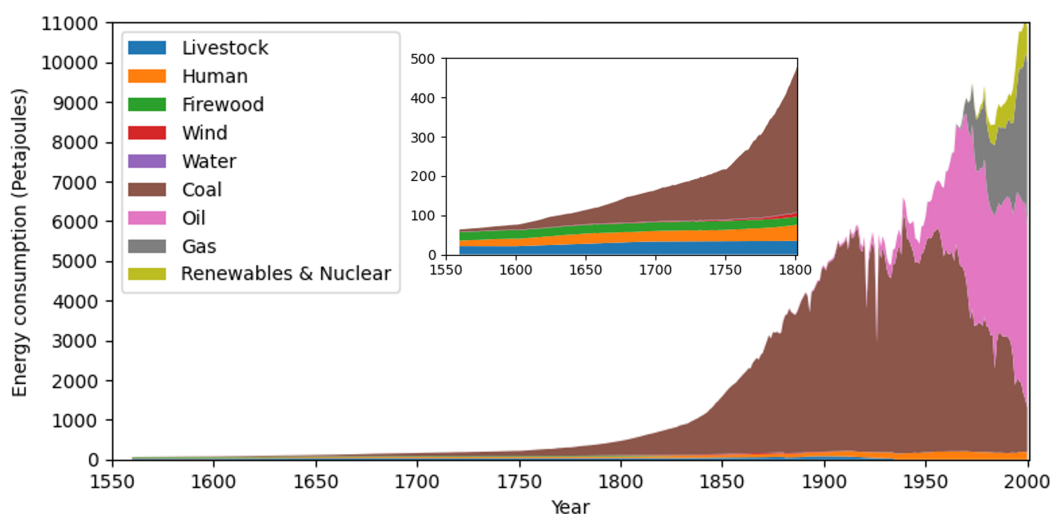


Fig. 1.1 Energy consumption in England and Wales by source. Insert shows consumption between 1560 and 1800. “Wind” and “Water” categories exclude wind and water used for generating electricity: these are included in the “Renewables & Nuclear” category [1].

on a current collector which can be directly connected to a circuit. Because LiCoO_2 has a higher reduction potential than graphite, the system is thermodynamically stable in this state. However, when the anode is connected to the negative terminal of a power source, and the cathode is connected to the positive terminal of a power source, electrons accumulate at the anode and diminish from the cathode. This charge imbalance makes it electrochemically favourable to reduce the anode and oxidise the cathode, and so Li^+ migrates from the cathode to the anode, balancing the charge difference. Eventually the system reaches equilibrium and the battery is considered to be fully charged. When the electrochemical cell is connected to a closed circuit without a power source (other than the battery), Li^+ migrate from LiC_6 to CoO_2 , which induces a flow of electrons through the circuit from the anode to the cathode. The current in the circuit transfers energy to the electrical components.


The choice of electrode materials is paramount to the performance of the battery. Batteries can only generate a current in a circuit so long as a charge carrier (usually Li^+) is flowing through the electrolyte to balance the charge, and therefore the amount of charge moving through the circuit is limited by the number of charge carriers that can be moved between

the electrodes. Consequently, electrodes are frequently discussed in terms of their “charge capacity” which is usually expressed in terms of the mass (gravimetric charge capacity) or its volume (volumetric charge capacity). The charge capacity is a measure of the number of charge carriers that can be delivered or accepted during the charge/discharge cycle. The theoretical charge capacity can usually be calculated directly but the practical charge capacity is usually slightly less. The difference between the theoretical and practical capacity can arise for a number of reasons, such as the capacity fading over the lifetime of the battery, but sometimes engineers impose constraints on the battery to improve its lifespan. The specific combination of electrodes is also important as the difference in reduction potential of the electrodes determines the potential difference across the cell, and the amount of energy that can be delivered by the battery when discharging is equal to the charge multiplied by the potential difference. With consideration to both electrodes, one can define an energy density for a battery (again as a gravimetric or volumetric property) which is determined by both the charge capacity and potential difference. Consideration should also be given to kinetics, as this will impact the discharge rate and therefore the power that can be delivered by the battery. Kinetics are influenced by numerous factors relating to the electrodes, the electrolyte, and a phase generated between the electrolyte and the electrodes called the solid-electrolyte interphase.

Despite considerable research into alternatives, the industry is still largely dominated by graphite/LiCoO₂ batteries [3, 4]. Many commercial LIBs (e.g. cobalt oxide, nickel manganese cobalt oxide, nickel cobalt aluminium oxide) achieve a specific energy density of 250-260 Wh kg⁻¹ [5], however, consumer’s desire for longer lasting devices is driving the development of electrodes with greater charge capacity. Conventional LIBs are made from intercalation electrodes: layered materials that allow for the insertion of Li⁺ between the layers when a suitable potential difference is applied. The disadvantage of this approach is the limited storage capacity for Li⁺ in intercalation materials. The conventional LIB

anode has a capacity of only one Li^+ ion for every six C atoms, limiting the theoretical charge capacity at 372 mAh g^{-1} [6]. Similarly, the conventional LIB cathode, LiCoO_2 , can intercalate one Li^+ ion for every three host atoms, which yields a theoretical capacity of 274 mAh g^{-1} . The strength of intercalation electrodes lies largely in their stability but they typically do not achieve a particularly high charge capacity. One possible alternative to intercalation electrodes are conversion electrodes [7–9].

Conversion electrodes perform the same role as intercalation electrodes but undergo a more complex phase evolution during the charging/discharging process. The addition of ions into an uncharged conversion electrode results in a disproportionation reaction that yields a mixed-phase electrode. A conversion electrode under a sufficiently large potential difference may undergo redox in the form,



where A and B are arbitrary elements and X , Y , and Z are arbitrary integers.

It is possible for conversion-type electrodes to achieve a much higher ratio of Li^+ ions [10] and therefore achieve a greater gravimetric capacity, such as Li_2S with a theoretical charge capacity of 1166 mAh/g [7]. But conversion-type electrodes come with a number of disadvantages. Ion kinetics have a significant impact on the rate of charge transfer in batteries and therefore a structure with open space to facilitate the movement of ions is usually desirable. Conversion electrodes feature significant structural changes and phase separation, which may require complex migrations of atoms and ions and which can lead to slow phase evolution. However, the greatest disadvantage of conversion electrodes is that these materials often undergo a significant change in volume during the charging/discharge process, which can lead to capacity fading.

The highest possible charge capacity for a LIB anode can be achieved by using pure Li metal and this approach has been explored since the 1970s. Successful commercialisation of

Li metal batteries (LMBs) are unfortunately seemingly out of reach due to safety concerns. Li metal undergoes a violent exothermic reaction in the presence of moisture and therefore there is a large hazard associated with damage to LMBs. Additionally, cycling LMBs leads to the formation of Li dendrites which can short circuit the battery, impeding its function and posing additional fire risk [11]. Alloy electrodes [12, 13, 11] are an emerging technology intended to approach the charge capacity of Li metal electrodes but with more stability and reduced dendrite formation. Like conversion electrodes, insertion and deinsertion of Li^+ leads to changes in the crystal structure and consequently they experience many of the same problems as conversion electrodes, such as capacity fading over time that is partially induced by the large volume changes. Insertion of Li^+ into alloy-electrodes leads to the formation of Li-alloys but the alloying material may be a metal such as Al, Sn, or Mg, or a non-metal like Si, P, or Ge [12].

Although LIBs are probably the most famous example of a electrochemical batteries, Li^+ is not the only possible charge carrier. Li^+ is appealing due to its low mass (facilitating low gravimetric capacity batteries) and its small ionic radius (facilitating good kinetics). However, Li's low abundance on earth and the difficulty of its extraction has made it expensive (further exacerbated by the high demand for batteries). In the search for alternatives, some have looked to using other alkali metals as charge carriers, particularly Na [14–16] for its wide accessibility, but also K [17, 18, 15, 16, 19]. K-ion batteries (KIBs) and Na-ion batteries (SIBs) offer many similar advantages, but while the great success of the graphite anode in LIBs over the past 30 years has made graphite an attractive choice for a new anode [18], the poor intercalation ratio of Na^+ into graphite layers [18] inherently leads to poor gravimetric/volumetric charge capacity. Na-graphite intercalation electrodes are therefore unlikely to be commercialised. K^+ is a promising alternative. It has been demonstrated that KIBs offer a lower reduction potential than SIBs in some organic solvents (leading to more efficient energy transfer), and as K^+ solvated ions are generally smaller than

Na^+ solvated ions, higher electrolyte conductivity and diffusion are expected[18]. Although graphite can only accommodate 75% as many K^+ ions as Li^+ ions per C atom, leading to an intrinsically lower charge capacity, the difference in accessibility between Li and K, may lead to a higher charge capacity per unit currency.

Non-alkali metal charge carriers, such as Mg^{2+} [12], S^{2-} [12], and Al^{3+} [20] have also been proposed but are less well studied. Much of the same theory concerning intercalation, conversion, and alloy electrodes also applies to alternative charge carriers.

Developing new batteries that are cheaper, longer-lived, and more energy dense necessitates the discovery of new materials. Unfortunately, there is no known method to directly find a material with a set of desired properties in the way that one might solve a mathematical equation. Consequently, there is a need to both screen known materials and *search* for new materials that might have desirable physical and chemical properties. For electrodes, these materials ideally should accommodate a large number of charge carriers per unit volume and mass; these properties make the batteries convenient to store and transport and when implemented in electric vehicles, also more energy efficient. The material should also have an open structure to facilitate fast charging and discharging, and a large potential difference between the anode and electrode are desirable as this affects the energy delivered by the battery via the relationship, $E = QV$,

$$E = QV, \quad (1.2)$$

where E is energy, Q is charge, and V is the potential difference. Finally, as with many materials, the electrode should be stable, cheap, and construction and disposal should have a minimal impact on the environment.

Structure searching has been performed experimentally by humans for millennia, with early examples taking the form of simple mixing of known substances in the hopes of discovering something new. Later developments in chemistry have lead to more sophisticated

approaches. In some cases, chemists identify a material with desirable features, and an understanding of reaction mechanisms and the periodic table allows for the design of reaction routes that may synthesise an analogous material with different elements. These efforts benefit greatly from the innovations in crystallography in the mid 20th-century which have facilitated accurate identification of crystal structures, but also the creation of online structure databases, such as the ICSD [21] and the CCDC [22], which have made crystal structures more widely accessible. Furthermore, innovations in industrialisation have allowed for high-throughput experimental material discovery through combinatorial chemistry in which large quantities of potential new materials are prepared automatically from a range of chemicals. However, synthesising new materials is sometimes expensive as synthetic routes may be complex and long, yields may be low, precursors may be expensive, and the proposed synthetic route may fail to yield the desired material. An alternative, cheaper approach is computational structure prediction, which was first made practical in the late 20th century due to advances in physics, physical chemistry, and statistical models, as well as ever increasing computing power.

Predicting whether a proposed structure is stable necessitates the prediction of a structure's total (or formation) energy and forces, as well as the total (or formation) energy of other possible structures. The position of each atom relative to the position of other atoms in the structure, and the elements involved play a significant role in determining the energy of a structure, therefore structure prediction usually requires modelling proposed structures at an atomic level. However, coarse grain models, which treat groups of atoms together as a single "entity" are sometimes sufficient. It is also of note that for accurate energy and property predictions, it is sometimes necessary to treat atoms and electrons separately.

The number of atoms that can be modelled in a practical time frame depends on the complexity of the model and the amount of computing power available. Classical models, such as Lennard-Jones, are far less computationally demanding than *ab initio* models

such as density-functional theory (DFT) and Hartree-Fock (HF), but typically give less accurate/reliable energy/force/stress predictions. Additionally, some phenomena, such as electronic band structures, cannot be predicted without modelling electrons, which increases the complexity of the model. However, advances in computing (RAM, clock speed, core counts, GPUs, etc) and improvements in the efficiency of models have made it possible to apply *ab initio* models to problems that previously would have been considered too computationally demanding. Where *ab initio* models are too expensive and no classical model is available that accurately predicts the desired properties, machine learning potentials (MLPs) may be a viable alternative. MLPs are highly parameterised models which are usually fit to data from experiment or *ab initio* models. Given the low computational cost of MLPs and that they have been demonstrated to provide more accurate energy and force predictions than classical models in many cases [23], I expect that the role of classical models will largely be replaced by MLPs in the coming years. Similarly to experimentalists, modellers have improved their own structure-searching efforts by assembling internationally-accessible databases of modelled structures. Two such examples are the Materials Project and OQMD, which perform high-throughput DFT geometry optimisations on thousands of structures. In some cases, properties beyond the total energy, such as band structures, are also reported for these structures. These databases have been particularly valuable recently as they include large quantities of structures (150,000+ on Materials Project and 1,000,000+ on OQMD) with energies all computed with the same DFT settings, making them useful for training MLPs. The algorithms for structure searching have also evolved over time. Markov-chain Monte Carlo and simulated annealing are closely related methods that have been used with great success since the early 1980s [24, 25] but have been partially replaced by newer methods such as nested sampling [26] in 2010, and *ab initio* random structure searching [27] in 2011. Each structure searching method has its own strength and weaknesses and must be selected carefully for the use case so as to ensure that as much relevant data as possible is acquired in

a timely manner. For example, nested sampling is ideal for constructing pressure-temperature (P-T) phase diagrams as it allows for precise determination for the temperature of a phase transition, but due to the computational cost of the method, it is inefficient at searching for ground-state structures. Conversely, AIRSS is highly efficient at searching for ground-state structures but provides no information about phase stability at finite-temperatures and therefore it is a poor choice for constructing P-T phase diagrams. Atomistic models and structure searching algorithms are vast and diverse, however, the models and utilised in this research are detailed in sections 2.2 and 2.3, a broad overview of structure searching methods is found in section 2.1, and details of the specific implementations of structure searching methods are found in the relevant chapters of each project.

Chapter 2

Methods

2.1 Structure Prediction

Computational atomistic modelling is a powerful method for learning about the physical world. In many cases, modelling provides a cheaper, faster, and safer alternative to empirical / practical approaches. However, the two approaches are regularly used as complementary methods for providing a more complete understanding of the research topic. One such example is a paper on modelling amorphous alumina, which was published by the Morris group and the Grey group in 2023 [28]. Amorphous structures, by definition, have no repeating crystal structure, which prohibits structure determination through experimental crystallographic methods such as X-ray diffraction. Computational structure prediction can also be challenging as the lack of a well-defined crystal means that a model cell should reflect an *average* atomic arrangement. This is typically achieved by generating a large cell but the cost of modelling such a large cell may be prohibitive. Harper *et al* surmounted this by generating amorphous alumina cells with various parameters, modelling their NMR spectra and XAS spectra, and comparing to experimental spectra. It was then possible to identify the parameters that lead to cells that best matched experiment [28].

A necessary precursor to modelling a system's chemical/physical properties is knowledge of its atomic structure, which itself requires knowledge of the system's stoichiometry. Sometimes a structure is known from prior research, or it can be determined experimentally (such as by X-ray diffraction), but when a characterisation is incomplete, ambiguous, or entirely absent, computational structure prediction may be the most time and cost efficient method of identifying that structure

2.1.1 Potential Energy Surfaces

A potential energy surface (PES) is a mathematical model relating the parameter space of an atomistic system, composed of all properties which describe the state and position of all particles in the system, to the potential energy of that system. PESs include a number of significant features; most notably, there are three types of stationary points: maxima, minima, and saddle points. Stationary points, which are points in phase space that have a gradient of zero, are significant in the context of gradient descent and quasi-Newton optimisation algorithms (see Section 2.1.2), as many of these algorithms are designed to minimise gradient, usually in hopes of discovering the global minimum (the lowest energy point on the PES). Maxima rarely pose a problem as they can only be accessed by moving *up* a gradient. Saddlepoints are features in which the gradient increases in some directions but decreases in other directions. For a PES with two geometric parameters, these saddlepoints are saddle-like in shape. These points can be problematic as they are accessible through gradient descent, but they do not represent a minimised energy. Minima are points in the PES where any small movement in atomic positions results in an increase in energy. The lowest minimum in the PES is the global minimum and represents the most stable structure for that system at 0 K. The remaining feature of note is the “basin”. Whereas the previous features are all points on the PES, the basin is an area of the PES centred around a minimum. The exact edge of a basin is difficult to define objectively, and a basin may contain a number of sub-basins. Each

point on the PES inside any given basin can be considered a variant of the same structure and the point at which an atomic configuration is entirely distinct from a minima is arbitrary. It is also of note that the number of points on a PES that represent an equivalent structure is directly related to the partition function, and therefore the thermodynamic accessibility of a structure is related not only to the depth of the basin but also to the width of the basin. The probability, p , of a state, i , being occupied is related to the energy of the state, E_i , the Boltzmann constant, k_B , the temperature of the system, T , and the canonical partition function, Z , by the Boltzmann distribution, which is expressed as:

$$p_i = \frac{1}{Z} \exp\left(-\frac{E_i}{k_B T}\right) \quad (2.1)$$

Consequently, a system at finite temperature may not necessarily occupy the state with the lowest possible energy. Larger temperatures lead to an increase in energy in the system which makes higher energy states accessible. The system is likely to take on the structure that has the largest sum probability of each of its equivalent states.

Let us consider the structure searching problem in the context of a simplified PES. For a system as simple as a 2-atom Lennard-Jones cluster (i.e. two identical and rotationally invariant atoms), the optimal structure can be found directly as the problem can be reduced to optimising a single variable: the distance between the atoms. The PES of this system may be represented fully with only 2 dimensions, where the first dimension represents the distance between atoms, and the second dimension represents the energy of the system. The entire shape of the PES is expressed clearly and concisely simply by computing the energy of a 2-atom Lennard-Jones cluster and varying the interatomic distance (Figure 2.1).

If a third atom is introduced, the complexity of the problem increases drastically. It is now necessary to optimise no fewer than three variables, by treating the three atoms as the vertices of a triangle on a 2D plane (note that that this model does not work if all three atoms are in a line). With this approach, one can describe the system exactly with two angles

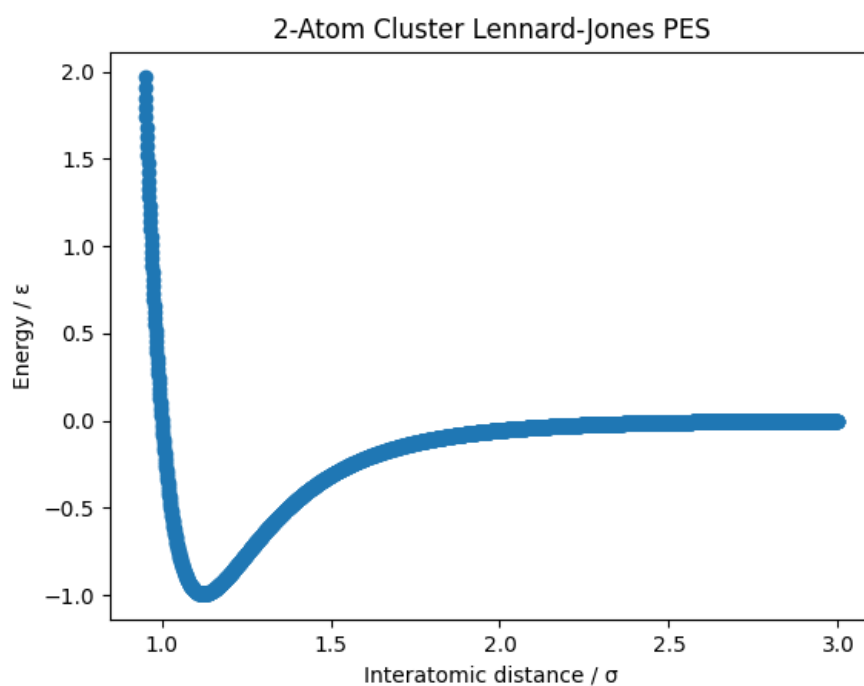


Fig. 2.1 Relationship between energy and interatomic distance for a 2-atom system. The exact shape of this 2D PES depends on the energy model used. Here, the Lennard-Jones pair-potential is used and has been parameterised such that $\epsilon = \sigma = 1$. Energy and distance are expressed in terms of ϵ and σ respectively, and the units are therefore dimensionless. The PES has been sampled with a step size of 0.01σ .

of the triangle, and the length of an edge between two vertices. The PES for this system would therefore be 4-dimensional, making it impractical to represent fully in a 3D world. To illustrate this, the PES for a 3-atom Lennard-Jones cluster has been included as Figure 2.2. In this plot, only the extremes of the PES are visible as three of the six faces of the cube. The majority of the data, which is contained inside the cube, is entirely obfuscated.

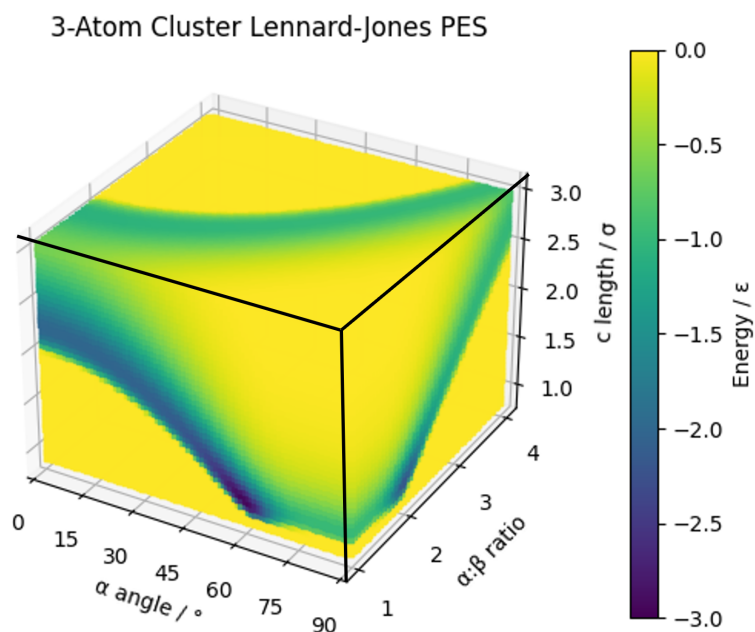


Fig. 2.2 Relationship between energy and geometry for a 3-atom system. The exact nature of this 4D PES depends on the energy model used. Here, the Lennard-Jones pair-potential is used and has been parameterised such that $\epsilon = \sigma = 1$. Energy and length are expressed in terms of ϵ and σ respectively, and the units are therefore dimensionless. The energy is shown in the plot with colour, with the minimal-energy geometry (an equilateral triangle with edge lengths of 1) shown in dark purple and higher energy regions shown in yellow. The Z-axis shows the length of the edge between two vertices with measured angles (c). The X-axis shows the angle of one of those vertices (α). Instead of showing the other angle, the Y-axis shows the ratio between the first angle and the second ($\alpha : \beta$) to avoid symmetric redundancy. The angle has been sampled with a step size of 1° . The angle ratio has been sampled with a step size of 0.05° . The edge length has been sampled with a step size of 0.05° .

A system with four atoms with P1 symmetry could be modelled as a 3D tetrahedron; a polyhedron constructed from four triangles. It would therefore be necessary to optimise no

fewer than nine variables (three distances and six angles). The complexity of the problem becomes progressively greater for larger systems.

So far only clusters of atoms in an infinite void have been considered, but when structure searching for materials, it is usually necessary to consider a periodic system with an effectively infinite number of atoms. The unit cell for a periodic system can be as small as a single atom but in order to know the periodicity of the system, cell parameters must be defined and to do this exactly. For a system with no symmetry, no fewer than 6 variables (3 lengths and 3 angles) are needed. Consequently, periodic systems have more variables to optimise than an equivalent non-periodic system. This is illustrated in Figure 2.3 for an orthorhombic cell (angles $\alpha = \beta = \gamma = 90^\circ$) containing only one atom. This should not be taken to mean that geometry optimisations are harder for periodic systems, however, as the complexity of the PES is not a simple function of the number of parameters.

It is therefore entirely impractical to attempt to minimise the energy of a system by tweaking geometric variables (for most systems). The search for the global minimum (the lowest energy structure for any given stoichiometry) is exacerbated by the fact that it is not usually possible to prove that a proposed structure is the lowest energy solution; it can only be proven that a structure is not the global minimum by discovering a lower energy structure. For this reason, a number of sophisticated approaches have been developed to navigate PESs.

2.1.2 Geometry Optimisations

The simplest example of a structure searching method is a geometry optimisation (see section 2.1.2 for further details), which is effective at locating the minimum of a basin but is not usually suitable for optimising between basins. Geometry optimisation algorithms have been designed to locate minima by altering atomic positions to minimise the magnitude of the

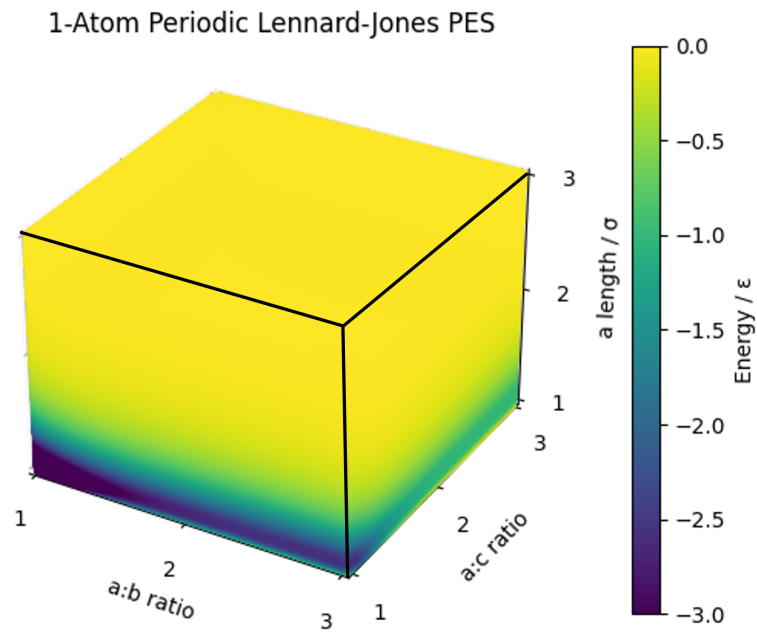


Fig. 2.3 Relationship between energy and geometry for a simple orthorhombic crystal. The exact nature of this 4D PES depends on the energy model used. Here, the Lennard-Jones pair-potential is used and has been parameterised such that $\epsilon = \sigma = 1$. Energy and length are expressed in terms of ϵ and σ respectively, and the units are therefore dimensionless. The energy is shown in the plot with colour, with the minimal-energy geometry shown in dark purple and higher energy regions shown in yellow. The Z-axis shows the length of the one of the edges of the cell (a). The X-axis and Y-axis show the ratio of the other two edges to the a edge to avoid symmetric redundancy. All variables have been sampled with a step size of 0.05.

forces in the system. The force on an atom, F , is related to the potential energy surface by

$$F = -\frac{dE}{dx} \quad (2.2)$$

where dE/dx is the change in total energy with respect to the distance moved by an atom, but also the gradient of the PES at a given point.

Consider, once again, two rotationally-invariant and identical atoms in infinite empty space. This system has a 1-dimensional PES, which may be described in three parts. At long range, the energy of the system trends towards zero. At short range, the energy of the system trends towards infinity. At some point in between, the energy of the system is minimised (and has a negative value). At the minimum, the gradient of the potential energy surface is zero and therefore the force experienced by the atoms is also zero (see Equation 2.2). At short range, the gradient of the PES is negative (reflecting that the resultant interaction is a repulsive force), while at long range, the gradient of the PES is positive (reflecting that the resultant interaction is an attractive force). Furthermore, at very short range, the force magnitude is very large, while at very long range the force is near to zero; reflecting the difference in energy change for a finite atomic movement at different regions of the PES. Hence, the sign of a force shows the direction of the resultant force acting on an atom, and the magnitude of the force shows the change in energy for a finite atomic movement. Neither the magnitude nor the sign may be used to determine *how far* an atom should move. For a simple system, one may attempt to gain insight into this using the derivative of the force (the second derivative of the energy), but for 3D n-body systems, an iterative approach with small atomic movements is the only practical approach. However, for a 2D or 3D-system, the relative magnitude of the force in each of the cardinal directions (along with the sign) can be used to calculate a force vector with its own arbitrary magnitude.

The geometry optimisation problem differs slightly for periodic systems. To avoid optimising the positions of infinite atoms, geometry optimisations are usually performed on

the smallest repeating unit of the crystal. Defining such a repeating unit requires not only the atomic positions but also cell parameters. Sometimes it is convenient to fix or constrain components of a structure, such as the coordinates of the atoms, or the parameters of the cell. The description of geometry optimisations given so far has assumed that cell parameters are fixed, but if they are not known, they should usually be optimised. This requires consideration of the stress on the cell. Stress, σ , is a force over an area, and the stress on a cell, which can be represented as a 3×3 tensor in the form,

$$\sigma = \begin{bmatrix} \sigma_{xx} & \sigma_{xy} & \sigma_{xz} \\ \sigma_{yx} & \sigma_{yy} & \sigma_{yz} \\ \sigma_{zx} & \sigma_{zy} & \sigma_{zz} \end{bmatrix}, \quad (2.3)$$

can be used to optimise cell parameters.

Fixing parameters of a structure simplifies the PES (by making certain regions of the PES inaccessible), which accelerates the geometry optimisation, and if the parameter can be fixed to a known “correct” value, there is an increased chance the geometry optimisation will lead to the global minimum. If the relative atomic positions do not need to be optimised, forces can be treated as if they are all zero. If the cell angles are known but cell lengths are not, all stress components except the diagonals (σ_{xx} , σ_{yy} , σ_{zz}), should be treated as zero. If there is a desire to optimise only the volume, the ratio of diagonal components of the stress tensor should be fixed and all other components should be treated as zero. Similar principles may be applied to optimise only a single component of the cell, such as the X-length or the Y-length. Other more complex controls are possible, such as moving only specific atoms, or fixing atoms to a plane, but they are less frequently used and beyond the scope of this writing.

One approach to geometry optimisations is gradient descent, which can be summarised succinctly as,

$$X_{n+1} = X_n + \gamma F_X, \quad (2.4)$$

where X is the X -coordinate of an atom, γ is scalar step size, F_X is the force acting on the atom in the X -axis, and n is an iteration count. Gradient descent is a parametric algorithm and γ must be chosen carefully. Smaller values require more iterations (and therefore time) to converge, but if the γ is too large, it is possible to “jump over” the global minimum. More modern approaches to gradient descent, such as ADAM (adaptive moment estimation) [29] use a large γ at the start of the optimisation and reduce the value over the course of the optimisation. ADAM is widely used for optimising neural networks (see Equation 2.3), while quasi-Newton methods are more popular for geometry optimisations.

Newton’s method is slightly more complicated than gradient descent in that it also incorporates the second derivative of the optimised function. To minimise the magnitude of the gradient in a problem with a single variable, X , Newton’s method can be expressed as,

$$X_{n+1} = X_n - \frac{f'(X)}{f''(X)}, \quad (2.5)$$

where n is an iteration count and $f'(X)$ and $f''(X)$ are the first and second derivatives of X .

Atoms in 3D, have an X , Y , and Z coordinate, which can be summarised as the vector, \mathbf{X} . For higher dimensional problems such as this, $f'(X)$ is now a vector of derivatives, which in this case is the negative of the force vector, \mathbf{F} , and $f''(X)$ is expanded into a Hessian matrix, \mathbf{H}_f .

For a geometry optimisation, Newton’s method is expressed as

$$\mathbf{X}_{n+1} = \mathbf{X}_n + \frac{\mathbf{F}}{\mathbf{H}_f} \quad (2.6)$$

Unfortunately, Newton's method is impractical to implement in many circumstances as the computational complexity of the method scales as $O(n^3)$. The difference in cost between gradient descent and Newton's method is entirely attributed to the cost of computing the Hessian matrix. For atomistic geometry optimisations, there is the added consideration that the problem scales with the size of the system as the Hessian has the dimensions $3N \times 3N$, where N is the number of atoms. Consequently, a number of quasi-Newton methods have been developed which use an approximation to the Hessian matrix, such as BFGS (Broyden–Fletcher–Goldfarb–Shanno) [30–33]. In BFGS, the Hessian matrix is never computed exactly but instead the algorithm begins by making a guess of the Hessian, \mathbf{B}_n . All subsequent iterations then improve on this guess via:

$$\mathbf{B}_{n+1} = \mathbf{B}_n + \frac{\mathbf{y}_n \mathbf{y}_n^T}{\mathbf{y}_n^T \mathbf{s}_n} - \frac{\mathbf{B}_n \mathbf{s}_n \mathbf{s}_n^T \mathbf{B}_n^T}{\mathbf{s}_n^T \mathbf{B}_n \mathbf{s}_n} \quad (2.7)$$

which is dependent on only the previous approximation of the Hessian matrix, the step vector, \mathbf{s}_n ,

$$\mathbf{s}_n = -\frac{\mathbf{F}_n}{\mathbf{B}_n}, \quad (2.8)$$

and the change in force between steps, \mathbf{y}_n ,

$$\mathbf{y}_n = \mathbf{F}_n - \mathbf{F}_{n+1}. \quad (2.9)$$

The method is otherwise analogous to the equation for Newton's method given in equation 2.6, with \mathbf{H}_f substituted for \mathbf{B}_n . To optimise the shape and size of the cell, not only the forces on each of the atoms are needed but also the stress tensor on the cell. One approach to optimising the cell parameters may be to treat stress optimisations and force optimisations separately. This could involve alternating between a force geometry optimisation walk and a stress geometry optimisation walk but this can be a lengthy process and may result in driving the structure into a local minimum, which may be difficult to escape when switching to

the other walk. Alternatively, one could alternate between a force optimisation step and a stress optimisation step, but quasi-Newton methods usually estimate the Hessian using the previous estimate and the last step, and alternating between force optimisation steps and stress optimisation steps may result in inconsistent Hessian estimates. Instead of this, one can consider the stresses and forces together as part of a single quasi-Newton step. This can be done by introducing a 3×3 deformation tensor, which is initiated as

$$\varepsilon = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad (2.10)$$

This tensor is used to perform transformations on the cell parameters, atomic coordinates, forces, and stresses. The prior two transformations are performed simply by taking the dot product of the matrices and the transpose of the deformation tensor, while the forces are transformed by taking the dot product of the forces and the deformation tensor. For stresses, one first needs to compute the virial stress, σ_V ,

$$\sigma_V = -V \left(\begin{bmatrix} \sigma_{xx} & \sigma_{xy} & \sigma_{xz} \\ \sigma_{yx} & \sigma_{yy} & \sigma_{yz} \\ \sigma_{zx} & \sigma_{zy} & \sigma_{zz} \end{bmatrix} + 3 \begin{bmatrix} P & 0 & 0 \\ 0 & P & 0 \\ 0 & 0 & P \end{bmatrix} \right) \quad (2.11)$$

where V is the volume and P is the pressure. The virial is then transformed by solving a linear matrix equation for the deformation tensor and the transpose of the virial, and transpose the result. Now that both the deformation tensor and the virial stress tensor have been acquired, the cell parameters can be optimised as if they were they were three “extra” atoms. In this model, each row of the deformation tensor acts as a the coordinates of one of the atoms, and each row of the virial stress tensor is treated as if it were the forces acting on each of the respective atoms. The three extra atoms can then be incorporated into the Hessian and

their “coordinates” (elements of the deformation tensor) can be optimised in the same way as the coordinates of the “real” atoms. The new stress tensor is then computed as simply the virial stress tensor divided by the negative volume of the cell [34]. Through use of a Hessian with dimensions $3(N+3) \times 3(N+3)$, where N is the number of atoms, it is possible to optimise all atomic positions and cell parameters as part of a single step, allowing for efficient optimisation of the cell.

The LBFGS, BFGSLineSearch, and LBFGSLineSearch variants of BFGS are also available. L stands for “Limited Memory” and the variant is used where RAM needs to be conserved, and the Line Search variant ensures that the energy and forces decrease monotonically. Geometry optimisations, though limited, form a vital component of other structure searching methods. A number of methods exist and though they generally solve the same problem, many of them have unique strengths that can be leveraged for more unique circumstances.

2.1.3 *Ab Initio* Random Structure Searching

Ab initio random structure searching (AIRSS) is a modern approach pioneered by Chris Pickard and Richard Needs in 2006 [35], but described in greater detail in 2011 [27]. Many structure searching algorithms, such as MCMC (described overleaf), attempt to locate the global minimum by introducing some mechanism that allows the structure to increase in energy and thereby move out of one basin and into another. In AIRSS, an approach is taken to search each basin individually by generating random structures *en masse* and then performing a normal geometry optimisation. Under ideal conditions, the random structure generation step will generate a structure near the top of each basin, and then the geometry optimiser will find the minimum of that basin.

AIRSS is an extremely simple method of discovering local and global minima but care should be taken to ensure the generation of sensible random starting structures for the method to be efficient. Firstly, if searching for a structure with a specific stoichiometry, it is necessary

to fix the ratio of elements in the cell to the desired stoichiometry. Including only a single formula unit of atoms in the starting cell will minimise the cost of the subsequent geometry optimisations, and also reduce the number of possible structures that can be generated. However, one formula unit may be insufficient to define the global minimum so it may be necessary to search over a range of formula units. Additionally, the majority of the phase space represents high-energy gas-like structures which may require lengthy geometry optimisations. Furthermore, if the atoms are too close, depending on the model, the energy of the system may be miscalculated. In DFT with a frozen core approximation, for example, the density of the core electrons will be uninfluenced by the distance to neighbouring nuclei, even if those nuclei are within the core radius. Experimental bond distances should be considered when choosing a minimum interatomic distance during cell generation, and the experimental densities of similar cells should be considered when imposing a maximum volume. It is also common to impose a space group or other symmetry constraints. By generating only high symmetry cells, the area of the phase space being explored is reduced and therefore it can be explored more quickly. Furthermore, most solids (amorphous materials and quasicrystals are noteworthy exceptions) contain some kind of symmetry and so it is unlikely that the global minimum will have P1 symmetry (no symmetry).

AIRSS may also be used to search for point defects, including interstitial defects, substitutional defects, and vacancies. This may be done by starting with a pristine (non-defected) cell, define an area within that cell to create defects, and then randomly add and remove atoms within the constraints set by the user. Defects often break the symmetry in a cell and so the user may often forgo symmetry constraints, which are often applied when searching for bulk structures and defects [36–41].

2.1.4 Other Structure Searching Algorithms

Markov-chain Monte Carlo (MCMC) [25] and simulated annealing (SA) [24] are two of the oldest structure searching methods and are closely related. In these methods, a structure is generated randomly and then the atomic positions and/or cell parameters are perturbed slightly (this is called a step). After a step is performed, a random number between 0 and 1 is generated to determine whether the step is accepted or rejected, given an acceptance criteria. This acceptance criteria is based on the difference in energy before and after the step, and in the case of SA, a controlled parameter called temperature (which acts as a model for temperature). Occasionally accepting a step which results in an increase in the energy of a structure allows for the structure to leave the basin and explore the PES more thoroughly. In SA, the temperature of the system starts high but is reduced over the course of the simulation: driving the system to lower energy structures. Most structure searching methods search for the lowest energy structure possible (the global minimum) as this structure will be the most stable at 0 K. However, SA can be configured with a minimum temperature greater than 0 K in order to search for the most thermodynamically accessible structure at a desired temperature.

Molecular Dynamics (MD), is not strictly a structure searching method but is instead used to simulate how a system changes with time. MD exists in a number of different ensembles in which different parameters are fixed (such as temperature or volume). In all cases, the method is intended to model the motion of atoms with time. This can be used to determine which structure is kinetically accessible (as total energy is not enough to know whether a structure will actually appear experimentally) or it can be used to generate amorphous structures.

Nested Sampling (NS) [26] is a multi-step process that is initiated by randomly generating a large quantity of structures. The next step is to remove the highest energy structure from the set of walkers. Then a structure is chosen at random, and cloned. The cloned structure is then

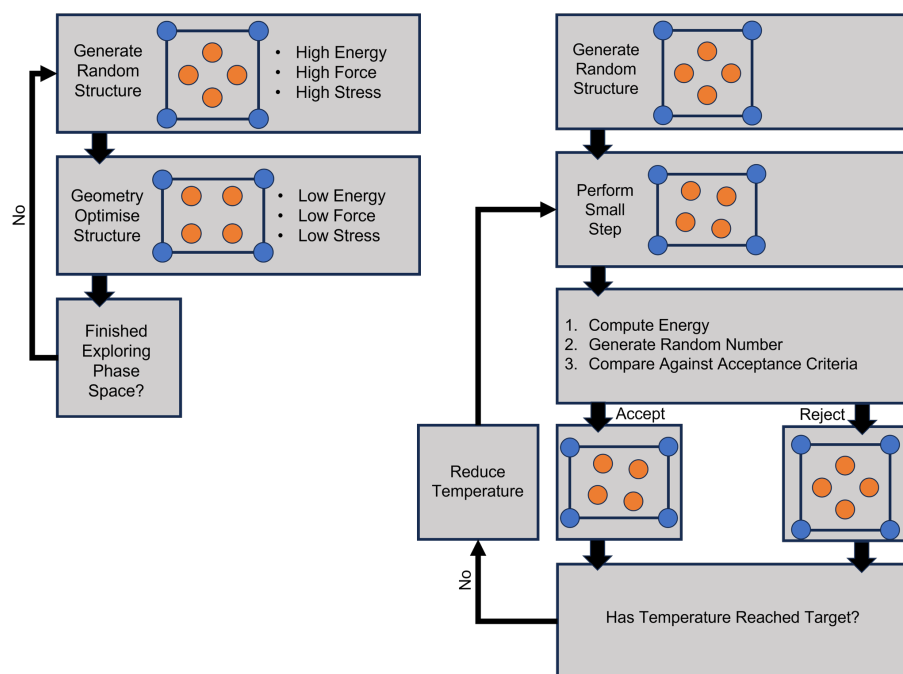


Fig. 2.4 Left: Algorithm diagram for AIRSS. Arbitrary atoms represented with blue and orange circles. Right: Algorithm diagram for Simulated Annealing. Arbitrary atoms represented with blue and orange circles.




Fig. 2.5 Molecular Dynamics example. Arbitrary atoms are represented with blue circles and direction and magnitude of vectors are given by red arrows and their lengths. Each cell represents a sequential time step.

subject to a MCMC step (if the MCMC step is rejected, further steps are attempted until one is accepted). The process then repeats from the point of removing a random structure until the total energy is minimised. NS is impractical for searching for a global minimum due to the large computational cost, but it is excellent at sampling a PES. Because the partition function can be found directly from NS, it is a powerful method for predicting phase transitions and therefore for constructing pressure-temperature phase diagrams.




Fig. 2.6 Nested Sampling Algorithm. Arbitrary atoms are represented with blue and orange circles.

Optimisation algorithms have also been proposed that were inspired by nature. Evolution is a natural and ongoing process by which species adapt to their environments and their competitors to maximise the likelihood of reproduction. The genetics of an organism are optimised by natural selection and genetic drift/mutation. Genetic algorithms (GA) are a

man-made approach to optimising other systems that are strongly based on these mechanisms. For atomistic structure searching applications, the algorithm begins by generating numerous random structures and then computing their energies (forces are not needed) [42–45]. Structures are then ranked according to their energy and each is assigned a random chance of being removed from the total pool of structures, with higher energy structures being more likely to be removed. The surviving structures are then randomly paired with another structure and “mated” to produce a new structure with influences from both parents. The children are often also subject to a small random perturbation (a mutation) to the atomic coordinates and/or cell parameters, which allows for the exploration of a wider phase space. The cycle of killing structures and reproducing is then repeated with subsequent generations of children until the algorithm converges.

Similarly, particle swarm optimisation (PSO) has been inspired by swarms of birds and insects. These swarms are composed of numerous individuals who follow others in the swarm and all (or most) converge on the same target. For atomistic structure searching applications, each candidate structure is considered a “particle” with “coordinates” comprised of many of the key variables of the structure (such as cell lengths and angles in the case of periodic systems). Each optimisable variable in the structure has a “velocity” associated with it, which acts as the step (in the language of optimisation algorithms) that each variable will take during the next optimisation iteration [46–48]. As with GAs, the algorithm is initialised by generating a number of random structures and then computing their energies. The algorithm then proceeds through an iterative process where each optimisable variable is modified by its velocity such that with each iteration, the coordinates of each particle become closer to the coordinates associated with the lowest known energy. After each iteration, the velocity of each variable is updated as a function of the current velocity, the lowest energy coordinates discovered by the given particle, the lowest energy coordinates discovered by any particle in the system, a set of user-defined parameters, and two variables that are randomised each

iteration. The algorithm proceeds until the particles converge or the algorithm reaches some maximum number of iterations.

GAs and PSOs have the key advantage that they do not require the differential of the optimised function. For atomistic structure searching, the force can usually be calculated but computing forces is often more expensive than only computing energies. However, PSOs and GAs in particular are highly parameterised and the success of the approach depends on selecting appropriate values for each parameter for each studied system. Additionally, as with NS, these algorithms do not treat the optimisation of each structure independently. The optimisation of each structure/walker is influenced by the other structures which are being optimised. Consequently, there is a positive relationship between the efficiency of each structure optimisation and then number of structures being optimised. However, in the case of AIRSS or MCMC/SA, the likelihood of any given structure reaching the global minimum is entirely independent of the number of structures being optimised.

The list of optimisation algorithms above is by no means comprehensive but many of the other algorithms operate under similar principals and processes.

2.2 Kohn-Sham DFT

Many modern condensed matter computer models can broadly be categorised into three sets: molecular mechanics (classical mechanics), *ab initio* (quantum mechanics), and machine learning (see Section 2.3). Molecular mechanics calculations boast a significant advantage in computational cost (performing force evaluations 2-3 orders of magnitude faster [49]), allowing for the modelling of much larger systems in a shorter period of time, but classical mechanics is entirely inadequate for describing certain electronic phenomena, such as electronic band structures. Machine learning approaches may be an appropriate compromise with a similar computational cost to classical methods and are *potentially* as accurate as *ab initio* methods. However, machine learning models are only as good as the data they are




Fig. 2.7 Left: Genetic Algorithm diagram. Arbitrary atoms represented with blue and orange circles. Individual structures are shown as stars with the difference between structures represented with colour. Right: Algorithm diagram for Particle Swarm Optimisation. Arbitrary atoms represented with blue and orange circles. Small block dots represent individual structures and direction and magnitude of vectors are given by red arrows and their lengths. The lowest energy structures is shown in green for a given iteration.

trained on, which may be from experiment or an *ab initio* model. There is therefore a need for accurate *ab initio* modelling.

The non-relativistic behaviour of electrons is described exactly by the many-body Schrödinger equation [50],

$$\hat{H}\psi = E\psi, \quad (2.12)$$

where, ψ is the wavefunction for a set of particles, E is the energy of the system, and \hat{H} is the Hamiltonian which for a system of nuclei and electrons can be expressed as,

$$\hat{H} = -\frac{\hbar^2}{2m_e} \sum_i \nabla_i^2 - \frac{\hbar^2}{2M_I} \sum_I \nabla_I^2 - \sum_{i,I} \frac{Z_I e^2}{|\mathbf{r}_i - \mathbf{R}_I|} + \frac{1}{2} \sum_{i \neq j} \frac{e^2}{|\mathbf{r}_i - \mathbf{r}_j|} + \frac{1}{2} \sum_{I \neq J} \frac{Z_I Z_J e^2}{|\mathbf{R}_I - \mathbf{R}_J|}, \quad (2.13)$$

where the five components represent the electron kinetic energy, nuclei kinetic energy, and electric potential energy associated with the following bodies: electron-nucleus, electron-electron, and nucleus-nucleus, respectively. These components are themselves expressed in terms of the reduced Planck constant, \hbar , elementary charge, e , the mass of an electron, m_e , nuclear mass, M , nuclear charge number, Z , electron coordinates, \mathbf{r} , and the nuclear coordinates, \mathbf{R} . In practice, this may be simplified via the Born-Oppenheimer approximation [51]. The mass of nucleons are three orders of magnitude greater than that of electrons and therefore nuclei travel more slowly. Hence, on a nuclear time scale, electron wavefunctions react to a change in nuclear positions almost instantaneously, which allows us to treat nuclear positions as if they are fixed. If the nuclei are immobile, the nuclear kinetic energy is 0, and nucleus-nucleus potential energy becomes a constant, and the system behaves as if it were at 0 K. One can therefore exclude these terms from the Hamiltonian.

Although the Born-Oppenheimer approximation simplifies the problem, the Hamiltonian above is still expressed in terms of multiple bodies and is extremely computationally demanding for more than a few particles. For this reason, it is rare to solve the electronic Schrödinger equation for a system more complex than a H atom. Because of the complexity of the many-body Schrödinger equation, condensed matter modelling is largely dominated by density-functional theory (DFT) which is built on two theorems [52, 53].

- The first Hohenberg-Kohn theorem states that the external potential, v_{ext} , is a unique functional of only the electron density, n .
- The second Hohenberg-Kohn theorem states that the electronic density which minimises the energy of the system is the true ground state.

Given the Hohenberg-Kohn theorems, the electron density that minimises the energy of the system is sufficient to compute all properties of a system at 0 K. However, computing the electron density is still a many-body problem and therefore an approximation is desirable.

In 1965, Walter Kohn and Lu Jeu Sham published Schrödinger equations for fictitious non-interacting electrons that generate the same electron density as interacting electrons [54]. The Kohn-Sham (KS) functional can be expressed in terms of three components: the kinetic energy of the non-interacting KS particles,

$$T_R[n] = -\frac{\hbar^2}{m_e} \sum_{i=1}^N \langle \varphi_i | \nabla^2 | \varphi_i \rangle, \quad (2.14)$$

(where φ is a single-particle wavefunction and N is the number of KS particles), the energy of an electron in a potential of its own charge density (the Hartree energy),

$$E_H[n] = \frac{1}{2} \int \int \frac{n(r)n(r')}{|r-r'|} dr dr', \quad (2.15)$$

and the exchange-correlation energy, E_{XC} , of an unknown form that contains both positive and negative energy contributions from electrons in a potential generated by *other* electrons, as well as the difference between the kinetic energy of interacting and non-interacting KS particles. The total energy of the system can therefore be expressed as:

$$E_{KS}[n] = T_R[n] + \int \rho(r) v_{ext}(r) dr + E_H[n] + E_{XC}[n], \quad (2.16)$$

where the additional component is the energy of an electron in an external potential.

KS DFT scales as $\mathcal{O}(n^3)$, which is much more efficient than the many-body Schrödinger equation, allowing for the calculation of the total energy of a simple many-atom system to be performed in a practical time frame on a home PC, while computing the total energy of the same system may be entirely impractical to solve on a Tier 1 HPC facility with the electronic Schrödinger equation.

2.2.1 Exchange-Correlation Functionals

While KS DFT is exact in principle, the exact form of $E_{XC}[n]$ is unknown and must be approximated. The exchange-correlation is often approximated with a functional, referred to as a “XC functional”.

The local density approximation (LDA) is a simple and widely used XC functional that was proposed in Kohn and Sham’s original paper [54]. This approximation assumes a homogeneous electron gas with an electron density that does not vary through space. With this approximation, the total XC energy can be expressed as an integral of the XC energy density over all space as

$$E_{XC}^{LDA}[n] = \int n(r) \epsilon_{XC}^{LDA}[n(r)] dr \quad (2.17)$$

$$\epsilon_{XC}^{LDA}[n(r)] = \frac{1}{2} \int \frac{n_{XC}^{LDA}(r, r')}{|r - r'|} dr' \quad (2.18)$$

$$n_{XC}^{LDA}(r, r') = n(r) \{g^h[|r - r'|, n(r)] - 1\}, \quad (2.19)$$

where g^h is the pair correlation function for a homogeneous electron gas.

Real systems do not have a homogeneous electron density and because correlation and exchange interactions occur between multiple electrons in different points of space, LDA, which poorly describes non-local effects, is an insufficient approximation in many cases. However, it has been widely used as a reasonable compromise between computational cost and accuracy for some systems. LDA performs particularly well for systems with relatively uniform electron density, such as bulk metals, but fails to reproduce properties of some other systems, such as the band gap of semiconductors and strongly-correlated materials, and ionisation energies in finite systems.

An improvement over LDA is the generalized gradient approximation (GGA), which includes a second-order gradient expansion with a term proportional to the squared gradient

of the density. GGA can be parameterised in a number of ways, and consequently, a large number of GGA functionals have been published. These can be parameterised either by fitting to experimental data or by deriving appropriate expressions from theoretical methods. GGAs better describe non-local effects and consequently give better predictions for bond lengths and angles, electronic band gaps, and other properties. Although GGA improves on these predictions, it is not exact and the accuracy of its predictions is inaccurate for some applications. Notoriously, both LDA and GGA underestimate the electron band gap due to a self-interaction effect, which delocalises the electronic density. Where more accurate electronic band gaps are required, such as when modelling optical properties, more accurate but computationally demanding approximations to the XC energy are usually applied. Examples include; XC functionals with higher order gradient expansions (meta-GGAs), DFT/Hartree-Fock hybrid approaches, and Hubbard U (which is used in combination with an XC functional). Due to their computational cost, these approaches are not conducive to high throughput structure searching. Consequently, they have not been used in this research and will not be detailed.

2.2.2 Semi-Empirical Dispersion Corrections

Dispersion interactions are long-range intermolecular interactions which arise from the coupling of electric fields at long range. In a classical context, this is often described as an "instantaneous dipole-induced dipole" interaction. Fluctuations in electron density can periodically form an instantaneous dipole in a molecule and the resulting electric field induces dipoles in neighbouring molecules, resulting in an attractive electrostatic force between them. While this model is useful for conceptualising dispersion interactions, the effect is much more accurately modelled by electron correlation; a phenomenon which is not described exactly in KS DFT.

Although local and semi-local functionals, such as LDA and GGA, are sufficient to describe the energy of many systems, dispersion interactions play a significant role in others; particularly layered systems (see section 5). To this end, semi-empirical dispersion corrections (SEDCs) have been developed to bridge the gap between KS-DFT and exact-correlation. The dispersion-corrected energy, E_{KS-D} , of a system made be expressed as,

$$E_{KS-D} = E_{KS} + E_{disp}, \quad (2.20)$$

where E_{KS} is the Kohn-Sham energy. E_{disp} depends on the implementation, but in the case of the G06 SEDC (which is the only SEDC used in this work), it takes the form,

$$E_{disp} = -s_6 \sum_{i=1}^{N_{at}-1} \sum_{j=i+1}^{N_{at}} \frac{C_6^{ij}}{R_6^{ij}} f_{dmp}(R_{ij}), \quad (2.21)$$

where, s_6 is a global scaling factor (which is fit empirically), N_{at} is the number of atoms in the cell, C_6^{ij} is a dispersion coefficient for atom pair ij (also fit empirically), R_{ij} is an interatomic distance, and f_{dmp} is a damping function of the form,

$$f_{dmp} = \frac{1}{1 + e^{-d(R_{ij}/R_r-1)}}, \quad (2.22)$$

to prevent near-singularities for very small values of R_{ij} . Here, R_r is the sum of atomic Van der Walls radii. [55]

A table of the DFT parameters for each project, which includes whether an SEDC was used, can be found in the supplemental data (section 9).

2.2.3 Pseudopotentials

The electrons in an atom may be separated into two categories: valence electrons, which have the largest principal quantum number, n , of any of the electrons associated with the

atom, and core electrons, which are the remaining electrons in the atom. The chemistry of atoms is dominated by the behaviour of valence electrons, because the core electrons are shielded from neighbouring atoms by the valence electrons. For structure searching, it is therefore less important to accurately model the core electrons than the valence electrons. Furthermore, KS wavefunctions, ϕ , display significant oscillations in proximity to the nucleus and representing these oscillations is computationally expensive with a plane wave basis set (see Section 2.2.4). Although all-electron DFT calculations have been performed [56], the cost of such calculations can be reduced by making an approximation to the core region of the wavefunction.

The electron density in most or all of the core region may be fixed such that it is independent of the chemical environment. This allows for the rapidly-varying region of the wavefunction near to the nucleus can be replaced with a pseudo-wavefunction with fewer nodes. The external potential is a unique functional of the electron density, which is the sum of the square of the KS wavefunctions, and so replacing a wavefunction with a pseudo-wavefunction also replaces the external potential with a pseudopotential (PSP). To this end, the pseudo-wavefunction must be constructed carefully such that the PSP smoothly joins with the all-electron potential beyond the core region. While many PSPs are constructed such that the core electron density is invariable with respect to the chemical environment, PSPs without the frozen core approximation are also available.

Two most common classes of PSPs are norm-conserving PSPs (NCP) and ultrasoft PSPs (USP). NCPs are constructed such that the logarithmic derivative of the all-electron wavefunction and pseudo-wavefunction vary in the same way. This ensures that for small changes in the external potential, the all-electron wavefunction and pseudo-wavefunctions are consistent within the valence region. NCPs are constructed such that the norm of the pseudo-wavefunction is identical to that of all-electron wavefunction within the core region. However, it is not strictly necessary for the norm of the all-electron wavefunction to equal

the pseudo-wavefunction. PSPs have therefore been developed that do not conserve the norm but are computationally cheaper to represent [57]. Because these PSPs are smooth and vary slowly, they are considered to be “ultra soft” and are therefore called “ultrasoft”. USPs delocalise charge from core orbitals into valence orbitals and are more accurate for atoms that do not have strongly bound core orbitals. USPs are complex to construct and may be technically challenging to implement; therefore, depending on the DFT package, USPs may not be supported for certain applications (for example, USPs are not supported for Raman or magnetic shielding calculations in CASTEP 21).

2.2.4 Bloch Theorem and Plane waves

Despite all the steps that have been taken to reduce the computational cost of *ab initio* calculations, modelling infinite crystals is still far beyond practical reach. Fortunately, it is usually only necessary to model the smallest repeating unit of the system, the unit cell. In a periodic system the electron density shares the same periodicity as the system itself and therefore so does the external potential (see Hohenberg-Kohn’s first theorem). For a periodic system, a wavefunction can be expressed, as in Bloch’s theorem, in terms of a periodic function, $u_k(r)$, and an arbitrary phase factor:

$$\psi_{\mathbf{k}}(r) = e^{i\mathbf{k}\cdot r} u_{\mathbf{k}}(\mathbf{r}) \quad (2.23)$$

and given that,

$$u_{\mathbf{k}}(\mathbf{r} + \mathbf{a}_i) = u_{\mathbf{k}}(\mathbf{r}), \quad (2.24)$$

where \mathbf{a}_i is the periodicity in a given axis, it follows that,

$$\psi_{\mathbf{k}}(\mathbf{r} + \mathbf{a}_i) = e^{i\mathbf{k}\cdot(\mathbf{r}+\mathbf{a}_i)} u_{\mathbf{k}}(\mathbf{r} + \mathbf{a}_i), \quad (2.25)$$

and hence,

$$\psi_{\mathbf{k}}(\mathbf{r} + \mathbf{a}_i) = e^{i\mathbf{k} \cdot \mathbf{a}_i} \psi_{\mathbf{k}}(\mathbf{r}). \quad (2.26)$$

Here, \mathbf{k} is the wave vector, which is the sole variable determining the phase of the wavefunction. A continuum of wave vectors are expressed in terms of a reciprocal space called k-space. k-space is itself periodic with a periodicity, \mathbf{b}_x , which depends on the real space lattice vectors, \mathbf{a}_x , in each axis, x as,

$$\mathbf{b}_1 = \frac{2\pi \cdot \mathbf{a}_2 \times \mathbf{a}_3}{|\mathbf{a}_1 \cdot \mathbf{a}_2 \times \mathbf{a}_3|}; \mathbf{b}_2 = \frac{2\pi \cdot \mathbf{a}_3 \times \mathbf{a}_1}{|\mathbf{a}_1 \cdot \mathbf{a}_2 \times \mathbf{a}_3|}; \mathbf{b}_3 = \frac{2\pi \cdot \mathbf{a}_1 \times \mathbf{a}_2}{|\mathbf{a}_1 \cdot \mathbf{a}_2 \times \mathbf{a}_3|} \quad (2.27)$$

The smallest repeating unit of k-space is called the first Brillouin Zone (BZ) and contains all unique wave vectors necessary for describing a wavefunction and hence the electron density. In principle, it is necessary to integrate over the whole BZ to determine the exact electron density, but because the energy of a wavefunction varies smoothly and continuously through most of k-space, it is a sufficient approximation to sum over a sample of weighted ($\omega_{\mathbf{k}}$) points in k-space (k-points) within the first BZ as [58],

$$n(\mathbf{r}) = \sum_j \sum_{\mathbf{k} \in \text{BZ}} \omega_{\mathbf{k}} |\phi_j^{\mathbf{k}}(\mathbf{r})|^2, \quad (2.28)$$

where j is an eigenstate. The sampled points k-points are usually, but not always homogeneously distributed as a Monkhorst-Pack grid [59, 60] throughout the first BZ or an irreducible representation (see below) of the first BZ. For any BZ with symmetry, it is unnecessary to sample the whole of the BZ. All unique features of the BZ are contained within an irreducible representation. In order to avoid computing “redundant” k-points and hence reduce computational cost, k-space may be sampled within this irreducible representation instead of the BZ as a whole.

A mathematical representation for one-electron orbitals is needed for Kohn-Sham DFT and for periodic systems, the basis set of preference is plane waves of the form,

$$e^{i\mathbf{G}\cdot\mathbf{r}}, \quad (2.29)$$

where \mathbf{G} is a reciprocal lattice vector. Usefully, a Fourier series of plane waves can be used to express a periodic function,

$$u_{\mathbf{k}}(r) = \sum_{\mathbf{G}=0}^{E^{cut}} C_{\mathbf{k}}(\mathbf{G}) e^{i\mathbf{G}\cdot\mathbf{r}}, \quad (2.30)$$

The periodic function is expressed as a Fourier series of an *infinite* number of plane waves but may be approximated to a *finite* number of plane waves. The kinetic energy associated with a plane wave, $T_{\mathbf{G}}^{\mathbf{k}}$, can be expressed in terms of the wave vector and a lattice vector via the relationship,

$$T_{\mathbf{G}}^{\mathbf{k}} = \frac{\hbar^2}{2m} |\mathbf{k} + \mathbf{G}|^2 \delta_{\mathbf{G}} \quad (2.31)$$

The Fourier coefficients, C_k , of progressively higher energy plane waves diminish in magnitude, and so a plane wave cutoff may be set. All reciprocal lattice vectors associated with a plane wave with an energy that exceeds the threshold are excluded from the calculation. This cutoff may be expressed as:

$$E^{cut} = \frac{\hbar^2}{2m} |\mathbf{G}|^2 \quad (2.32)$$

Given equation 2.23 and equation 2.30, an approximation for a single-particle wave function, φ can be expressed, as [61]

$$\varphi_j^{\mathbf{k}}(\mathbf{r}) = e^{i\mathbf{k}\cdot\mathbf{r}} \sum_{\mathbf{G}=0}^{E^{cut}} C_{j\mathbf{k}}(\mathbf{G}) \phi_{\mathbf{G}}(\mathbf{r}) \quad (2.33)$$

where

$$\phi_{\mathbf{G}}(\mathbf{r}) = \frac{e^{i\mathbf{k}\cdot\mathbf{r}}}{\sqrt{\Omega}} \quad (2.34)$$

is the plane wave basis set and Ω is the reciprocal cell volume.

2.2.5 DFT in Practice

Thus far, how the energy of a system may be found from the ground-state electron density of that system has been described, but no introduction has been given for a practical means of identifying that ground state density. This is usually acquired iteratively (though ML methods of guessing the ground-state density are in development [62–67]) through the self-consistent field (SCF) method by guessing an initial starting density and using this to solve the KS equations, which returns a new density. That new density is again used to solve the KS equations and the process is repeated until the difference between the input and output density are within practical tolerances. The Schrödinger equation for a KS non-interacting particle, which is solved as part of the SCF cycle, may be written as:

$$\left(-\frac{1}{2}\nabla^2 + v_{KS}\right)\phi = \epsilon_i\phi(r), \quad (2.35)$$

in which the first term is the kinetic energy of the particle, v_{KS} is the KS potential,

$$v_{KS} = v_{ext}(r) + \int \frac{n_{in}(r')}{|r-r'|}dr' + \frac{\delta E_{XC}[n]}{\delta n(r)}, \quad (2.36)$$

and ϵ_i is an energy eigenstate.

SCF cycles suffer from large charge displacements between iterations, which leads to rapidly changing v_{KS} and poor convergence. This effect can be damped by mixing the input and output densities. Several mixing schemes are available, but all work in this thesis utilised

the Pulay mixing scheme [68], which may be expressed as,

$$n_{in}^{x+1}(r) = \alpha \bar{n}_{out}^x(r) + (1 - \alpha) \bar{x}_{in}^n(r), \quad (2.37)$$

$$\bar{n}_{in}^x(r) = \sum_{i=1}^P \beta_i n_{in}^{(x-P+1)}(r), \quad (2.38)$$

$$\bar{n}_{out}^x(r) = \sum_{i=1}^P \beta_i n_{out}^{(x-P+1)}(r), \quad (2.39)$$

where P is the number of iterations (typically between 3 and 5), x , is the iteration, and α and β are coefficients with values between 0 and 1.

An SCF cycle for KS DFT with a plane wave basis set may be summarised as:

1. Guess electron density
2. Compute KS potential from electron density
3. Diagonalise KS potential for plane wave basis set, as in equation 2.34
4. Compute single-particle wavefunction from plane wave basis set, as in equation 2.33
5. Compute total electron density, as in equation 2.28
6. Modify electron density with density-mixing, as in equation 2.37
7. Compute total energy of system, as in equations 2.16
8. If difference between energy of this iteration and the previous iteration are greater than tolerance, repeat from step 2, otherwise complete SCF cycle.

Unless specified otherwise, all DFT calculations performed as part of this research were done so using CASTEP [69]. CASTEP is a plane wave Kohn-Sham DFT package implemented in Fortran. CASTEP supports a broad range of PSPs and XC functionals, as well as numerous different modelling methods such as finite-difference phonon modelling

and magnetic shielding calculations. Specific details of the parameters used for each project can be found in the Supplemental Data (Section 9).

2.3 Artificial Neural Networks

Artificial Neural Networks (ANN) date back to the first half of the 20th century [70], however they did not see widespread success or popularity until the last decade. ANNs are highly complex and flexible constructs that are capable of approximating complex functions to a high degree of accuracy (in some cases).

2.3.1 Early Architecture

Early models were called ‘Perceptrons’ and were based on the contemporary understanding of animal brains. Perceptrons consist of three types of layers, an input layer, a hidden layer (early perceptrons usually only had one hidden layer but modern perceptrons normally have multiple), and an output layer. Each layer contains a number of nodes (which are thought to mimic neurons) and each of these nodes has a unidirectional connection to each node in the neighbouring layers. The input nodes accept any real numerical value which is propagated through the hidden layers and eventually into the output nodes. Simple perceptrons have two types of parameters that can be optimised: weights, and biases. It is these parameters that give each hidden node its flexibility, and by extension, the ability for a neural network to learn a function. Each hidden node can be expressed as,

$$x = B + \sum_{i=1}^{i=n} A_i W_i \quad (2.40)$$

where B is the bias of a neuron, W is the weight of each input connection, n is the number of input connections, A is the input of the connection, and x is the output of the node. Not all nodes in a network are equally important for an accurate output, but until a network is trained,

the importance of each node is unknown. Activation functions have been implemented as an automated means of “activating” important nodes and “deactivating” unimportant nodes by, generally speaking, making large values larger and small numbers smaller, or zero. The simplest activation function is the identity, in which $f(x) = x$. A more widely used example is the sigmoid activation function: [71]

$$f(x) = \frac{1}{1 + e^{-x}} \quad (2.41)$$

2.3.2 Training

ANNs are understood to “learn” a function, and to this end they must be “trained” to minimise a loss function. The loss functions serves to measure the error in the ANN and can come in numerous different forms. In this work, the loss is usually based on the mean square error of the prediction. A loss function for energy predictions may be expressed as,

$$Loss = MSE = \frac{1}{n} \sum_{i=1}^n (E_i - \hat{E}_i)^2, \quad (2.42)$$

where E_i is the reference energy and \hat{E}_i is the energy prediction. Sections 2.1.1 and 2.1.2 gave an introduction to potential energy surfaces and optimisation algorithms. These principles can in large part also be applied to training ANNs though the algorithms that work well for geometry optimisations are not necessarily the same algorithms that work well for ANNs. Gradient descent and modern relatives are the preferred algorithms for optimising (training) the weights and biases of a neural network and in order to get the gradients, backpropagation is applied. Backpropagation is an application of the chain rule that finds the partial derivative of the loss with respect to the partial derivative of each weight and bias. Each of these partial

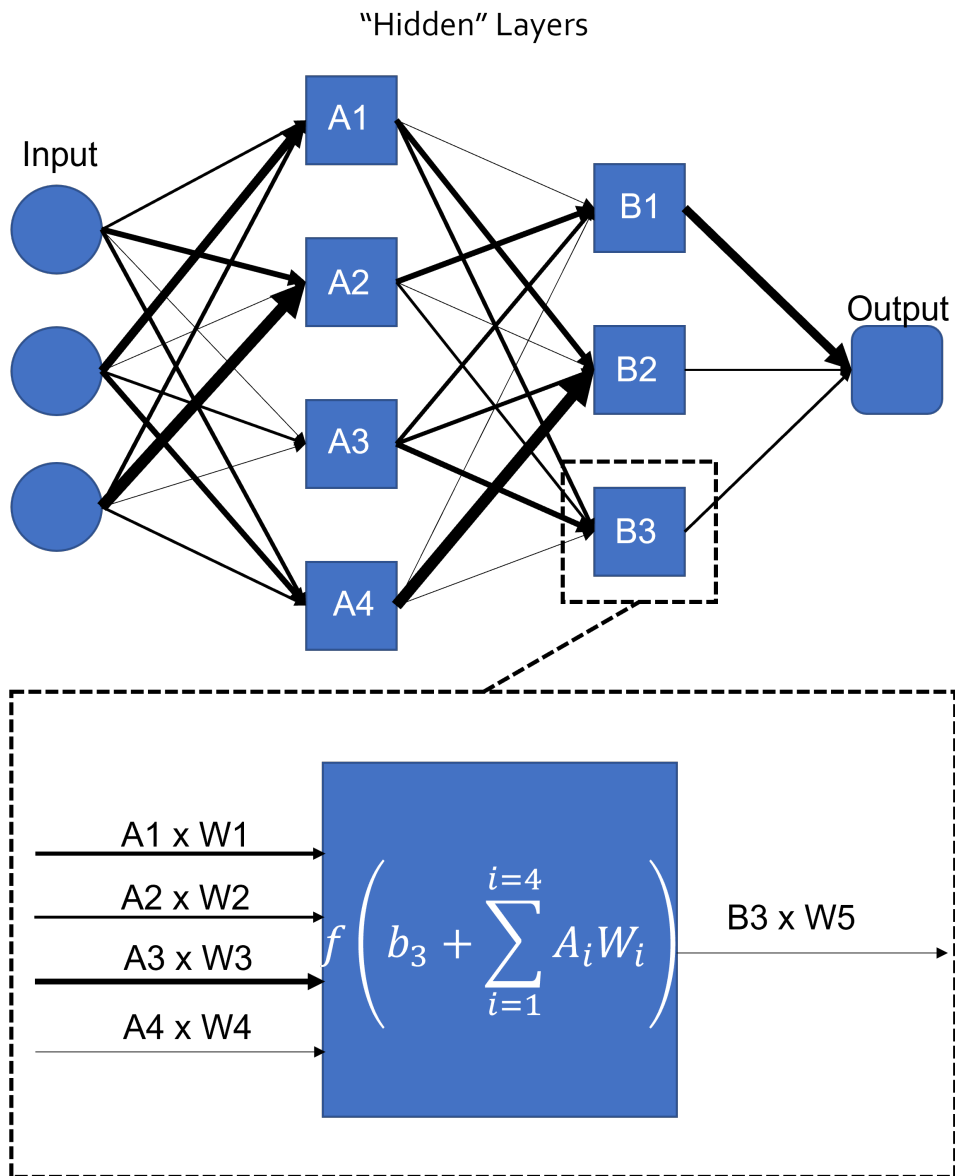


Fig. 2.8 Top: Architecture of fully-connected perceptron with two hidden layers. Outputs of first and second hidden layers are labelled with A and B , respectively. Bottom: Details of node $B3$ with arbitrary activation function, f . Weights are labelled with W and bias with b .

derivatives is calculated using partial derivatives of variables *later* in the ANN and so the gradients are thought to backpropagate through the network.

Consider a neural network with one input node, x , a connection, w_1 , a hidden node, h , another connection, w_2 , and an output node, y . For simplicity, the model NN features no biases or activation functions. The model also use the loss function defined in equation 2.42. The partial derivative of the loss, L , with respect to the partial derivative of w_2 is dependent only on themselves and the output of the ANN y , as,

$$\frac{\partial L}{\partial w_2} = \frac{\partial L}{\partial y} \frac{\partial y}{\partial w_2} = h(E - \hat{E}) = xw_1(E - \hat{E}), \quad (2.43)$$

while the analogous gradient for the first weight considers the loss, w_1 , y , and the output of the hidden node, h , as,

$$\frac{\partial L}{\partial w_1} = \frac{\partial L}{\partial y} \frac{\partial y}{\partial h} \frac{\partial h}{\partial w_1} = xw_2(E - \hat{E}), \quad (2.44)$$

where E is the energy prediction and \hat{E} is the reference energy. One can therefore initialise a neural network with random (but sensible) starting weights and biases, calculate the loss for a set of training data, and then use backpropagation to find the gradients for each parameter. An optimisation algorithm, such as ADAM, can then use these gradients to change the parameters to more favourable values. Then the loss is calculated again, gradients are calculated via backpropagation, the optimiser updates the parameters, and the process repeats until the loss converges [71].

Consideration must also be given to the training data itself. The training data should be diverse such that the ANN is asked to extrapolate as rarely as possible, but also, the number of “repeat” data points should be reflective of the frequency of occurrences of that data type in the set of data that the ANN will be applied against. For example, when training a NN that identifies dog breeds from photos, if half of all dogs are chihuahuas (they are not), half of the training set should be photos of chihuahuas. Additionally, it is generally true that the

larger the training set, the more accurate the ANN will be. However, it is not necessary to find the loss for every training point before updating the NNs parameters, and except for very small training sets, it is rarely done. In practice, the training set is broken up into smaller subsets called batches. Each iteration of the training process is performed on a different batch until all batches have been trained against, this is called an epoch. After every epoch, a new set of batches is generated randomly. Generally speaking, smaller batches lead to faster convergence but a less accurate ANN. Smaller batches also have the advantage of consuming less RAM.

A final point to consider with training is overfitting. For an ANN to be useful, it should be accurate but also generalisable. On initialisation of the ANN, the loss for all possible data points is relatively high and as the training progresses, this value will decrease, but it decreases faster for data points inside the training set than for data points outside the training set. If training progresses far enough, the loss for data outside the training set will eventually begin to increase. The ANN is overfitting. The most common way to mitigate this is by constructing a validation set of data. The validation set should be much smaller than the training set and the contents of the two sets should be mutually exclusive. After each batch (or epoch), the loss of the validation set should be calculated. When the validation loss begins to consistently increase, the model is overfitting. The “best” ANN, is the one that has the lowest validation loss. Overfitting can also be reduced by minimising the complexity of the architecture and by reducing the batch size.

2.3.3 Atomistic ANNs

Modern advances in computer hardware have facilitated the training of ANNs with much more complex architectures; most notably, the use of more than one hidden layer (deep ANNs). The accessibility of modern computer hardware, big data, and deep learning have all played a major role in the machine learning “renaissance” that has been occurring over

the last decade. Perceptrons are no longer the only architecture available; a plethora of new architectures have been developed and tailored to solve specific problems in different fields. Three of these architectures have been utilised for this research. Perceptrons (see above), convolutional neural networks (which are widely used for image processing tasks), and message passing neural networks (which are used for learning graphs, such as social connections).

For most problems, data in its raw format is not well suited as input for ANNs and atomistic systems are no exception. A number of different formats are available for storing information about atomic positions but generally the cell is stored either as three lengths and three angles, or as a 3×3 array, and the atomic positions are stored as either scaled positions or absolute positions with an ID (symbol or number) for the element. One could imagine a perceptron with six inputs for the cell information and then another four (an element ID and three coordinates) for each atom but because the number of inputs scales with the number of atoms, any trained ANN would only be applicable on systems with the same number of atoms. The problem of predicting the energy of a structure is therefore usually reduced to predicting the energy of an individual atom and then summing the values of all atoms in the cell to get the total energy of the structure. This approach would allow us to use only four inputs (or three if one ANN is wanted per element) to train an ANN that could be applied to any system, but regardless of whether scaled or absolute coordinates are used, there is no information about the shape or size of the cell and there is no information about neighbouring atoms. Consequently, a C atom in an infinite void is entirely indistinguishable from a C atom in a sheet of graphene. Any input used should contain information about the local chemical environment and it should also be translationally invariant and rotationally invariant (or equivariant) such that two identical cells in different orientations return the same results. One of the earliest successful basis functions designed to satisfy these conditions are (weighted) atomic centred symmetry functions, which work well but are relatively complex [72, 73].

More recently, success has been had by using much simpler radial basis functions convoluted through a range of trained filters, each of which provide a different representation for the chemical environment [74–76].

2.3.4 ANNs in this Project

Much of this research was carried out using the SchNet [74, 75] architecture (a continuous-filter convolutional neural network) as implemented in SchNetPack [77, 78], but this was later replaced by PaiNN [76] (an equivariant message passing neural network), also implemented in SchNetPack because it was demonstrated to give more accurate force and stress predictions. Because of this, all data presented in this thesis is the product of the PaiNN architecture.

The PaiNN architecture is intrinsically different to its predecessor, SchNet, and significantly more complex than the perceptrons described in section 2.3.1. The motivation behind the change in architecture is to reduce the size of the model and the training time, as well as improving the accuracy of property predictions; in particular tensorial properties such as stress [76].

Perceptrons and their closely related architectures (such as SchNet) are unidirectional, that is to say, data is transferred from one layer of nodes to the next, but never within a layer or back to the previous layer. In such an architecture, it is necessary to perform some kind of transformation to convert the information about each atom into a feature vector that forms the input. Alternatively, one can construct a graph neural network, in which each atom is represented by its own node, and each node is connected to another node if it is within some cutoff (neighbour) distance. The edges in this graph therefore mimic chemical bonds. Each node has a starting feature vector called an embedding, which in the case of PaiNN, is learned from a single perceptron layer (henceforth referred to as an “atomwise” layer). The value of each node is then passed to each immediately connected node via a “message” function, and the value of all nodes is then updated with the new value via an

“update function”. The message and update steps can be repeated as many times as desirable, propagating information through the network and eventually converging. This is called a message passing neural network. The message and update functions can include many of the architectural features seen in other ANNs, such as atomwise/perceptron layers, activation functions, and convolutions (all of which appear in PaiNN). With PaiNN, there are three message and update steps, followed by an atomwise layer, an activation function, another atomwise layer, and finally the output node, which is an energy prediction.

Although many atomistic ANN have implemented invariant atomic representations, Schütt *et al* raised that equivariant representations can resolve more structures [76]. Forces are vectors and stresses are tensors and each of their components have both a magnitude and a sign; given the importance of the sign, an equivariant representation therefore seems a sensible choice. The remaining property of concern is energy, which is a scalar. Energies are sometimes expressed with a negative or positive sign but unlike forces and stresses, this sign does not relate to a direction, but a difference in energy as compared to some reference point; invariance therefore seems sufficient to represent this property. Testing with PaiNN shows that errors in energy predictions are suitable, although greater than they were when trained on the SchNet architecture, which uses exclusively invariant representations. The PaiNN architecture implements two different continuous-filter convolutions as its representation. Convolutions have been utilised in SchNet and PaiNN because they are effective at highlighting geometric features. In the context of neural networks and in image processing, an array of data is convolved with a filter array to produce a new array of convolved data. Each element of the new array takes the value of the sum of the products of one element in each array such that a convolution between two 1D arrays, a and b , may be expressed as:

$$(a * b)_n = \sum_{i,j,i+j=n} a_i \cdot b_{j-1}, \quad (2.45)$$

where n is the index in the new array and i and j are indices in arrays a and b , respectively. Filters may either be learned by the ANN or designed in order to exaggerate certain features, such as in Figure 2.9, where a binary array representing a cross has been convoluted with a filter designed to exaggerate the upper edge of the cross. In the PaiNN architecture, the first




Fig. 2.9 Convolution between a binary array representing a cross (left), and a filter (center) designed to exaggerate the upper edge of the cross in the output array (right). Here, the edges of the output array have been cropped such that every element of the output array represents an operation in which the filter fits entirely inside the input array.

convolution is inherited from the SchNet architecture where the rotationally-invariant filters, W_s , are linear combinations of radial basis functions,

$$\sin\left(\frac{n\pi}{r_{cut}}\|\mathbf{r}_{ij}\|\right)/\|\mathbf{r}_{ij}\| \quad 1 \leq n \leq 20 \quad (2.46)$$

and the convolutions are performed on atomwise layers that as input, take embeddings learned from nuclear charges.

To aid with the training and implementation of MLPs, extensive use has been made of ML-Tools, a Python3 package developed as part of this research for the purpose of interfacing CASTEP with SchNetPack as well as providing additional tools that assist the user from the beginning to the end of the workflow. ML-Tools firstly provides the essential role of converting each geometry optimisation step from CASTEP files into ASE [79] atoms objects (which are widely used for storing cell data), and then storing that information in a database file that can be read by SchNetPack. The package also utilises the database's

metadata to store which structures belong to each geometry optimisation. Because of this, instead of relying on SchNetPack to assign structures to a training set or validation set randomly, a custom Python 3 script was used to assign structures to a set semi-randomly but in such a way that two structures from a single geometry optimisation are not assigned to different sets. This can be important as the steps from a geometry optimisation are highly correlated, particularly towards the end of an optimisation. If geometry optimisation walks are not limited to a single set, it is possible for structures which are effectively identical to be split between the training set and the validation/test set. Using a custom script for splitting structures between the training and validation set is also useful for cross validation (which is useful for hyperparameter optimisation), and for query by committee. ML-Tools also allows me to plot the convergence of the training and validation loss, which is useful for determining whether an ANN should be trained further, and for plotting MLP predicted properties against the DFT reference properties, providing insight about the accuracy of the model. ML-Tools also provides the means for performing singlepoint calculations, geometry optimisations, and molecular dynamics simulations via ASE [79] and using a SchNetPack MLP.

2.4 Observables

DFT and MLPs, discussed in the prior sections, are both suitable methods of modelling energies, forces, and stresses for atomistic systems. They are therefore frequently used for optimising structures and for comparing the stability of competing geometries, which makes them both a valuable tool in structure searching. However, once a structure has been predicted computationally, it cannot be directly compared to experimental data to verify the prediction. A wide range of analytical methods are employed by experimentalists to characterise unknown substances and in some cases these methods can be modelled. The ease of modelling an analytical method depends entirely on the natural phenomenon that makes the method possible. Infrared (IR) and Raman spectroscopy are analytical methods

of identifying chemical features in a sample through detection of photons passing through (in the case of IR spectroscopy) or scattered from (in the case of Raman spectroscopy). The energy of the detected photons differs from the energy of the incident photons due to the creation and absorption of phonons that result when the material interacts these photons. Any model with sufficiently accurate force predictions can be used to compute a phonon spectrum at the Γ point, and IR and Raman selectivity can be determined through space group theory. Therefore it is possible to model IR and Raman spectra with DFT or an MLP (though MLP force predictions are not usually good enough to do this well at present). However, NMR spectroscopy exploits the attenuation of magnetic fields by electrons. MLPs such as the ones used in this research do not model electron density and therefore they are unsuitable for such modelling, while DFT is frequently used for modelling NMR. However, due to the extreme flexibility of ANNs [80], it is expected that all known characterisation methods could be modelled by a specialised ANN. In this work, both phonons and NMR have been modelled using DFT.

2.4.1 Phonons

All atoms in molecules and crystals are vibrating constantly; even at 0 K. In molecules, there are $3N - 5$ degrees of freedom for linear molecules and $3N - 6$ degrees of freedom for nonlinear molecules, where N is the number of atoms. These degrees of freedom map directly to the number of possible vibrational modes and each mode has a discrete frequency (and energy). The frequency of light that can be absorbed to create such a phonon, which is conventionally expressed as a wavenumber, $\bar{\nu}$, and can be calculated directly from the reduced mass, μ , the spring constant k , and the speed of light, c , via the relationship:

$$\bar{\nu} = \frac{1}{2\pi c} \sqrt{\frac{k}{\mu}} \quad (2.47)$$

For crystals, which may contain an unlimited number of atoms, the question of “how many vibrational modes are possible?” is further complicated. In a crystal, atoms do not vibrate independently but instead occur as collective excitations called phonons. For a 1D chain, the change in atomic coordinates with respect to the equilibrium position, δx_n , can be expressed as a wave as

$$\delta x_n = A e^{i\omega t - i q n a_i} \quad (2.48)$$

in terms of the amplitude of the oscillation, A , the phonon frequency ω , the wave vector, q , an integer, n , and the periodicity, a_i .

This wave vector, q , is directly analogous to the wave vector, k , described in Section 2.2.4 but conventionally q is used for phonons and k is used for electron wave functions. Just as in Section 2.2.4, all unique phonon modes can be described in terms of the wave vectors in the first BZ. Consequently, any vibration with wave vector q is equivalent to any wave vector $q + b_i$, where b_i is the length of the BZ in a given axis. Although the wave vector was allowed to take any value within the first BZ in Section 2.2.4, it would be a mistake to do the same here. As there are a finite number of atoms in the unit cell, and hence a finite number of unit cells, a finite number of wave vectors are possible. For a 1D system, the wave vectors satisfy the condition,

$$q = \frac{2\pi p}{a_x}, \quad (2.49)$$

or for a supercell containing N unit cells,

$$q = \frac{2\pi p}{N a_x}, \quad (2.50)$$

where p is an integer and a_x is the lattice vector. From Section 2.2.4 it is known that the BZ contains all unique wave vectors and so only wave vectors within the BZ need to be considered for a full understanding of the phonon behaviour. The total number of normal modes in the crystal per degree of freedom can therefore be expressed in terms of the range

of q in the BZ, which for a 1D system is $2\pi/a_x$, and the spacing between allowed values of q , which is $2\pi/Na_x$ [81]

$$\text{Number of normal modes} = \frac{2\pi/a_x}{2\pi/Na_x} = N \quad (2.51)$$

It would appear therefore that a single unit cell (for which $N = 1$) only a single normal mode / wave vector is defined, which is the Γ point at the centre of the BZ. Each of these normal modes has a phonon frequency, ω , which depends on the geometry of the crystal, the mass of the particles, the strength of the chemical bonds, and the lattice wave vector, q . In a simplistic monatomic 1-D chain, where there is only one normal mode, the phonon frequency, ω , may be expressed as,

$$\omega = 2\sqrt{\frac{\kappa}{m}} \left| \sin \frac{qa}{2} \right|, \quad (2.52)$$

Where m is the mass of each atom, κ is the spring constant (determined by the chemical bonding), q is the lattice vector, and a is the lattice vector.

At the time of writing, CASTEP implements two method for computing phonons. The finite displacement method involves displacing each atom by a small quantity in each direction, and evaluating the forces on the perturbed configuration such that the change in energy, E , with respect to displacement, u , may be expressed as

$$\frac{d^2E}{du_{\kappa,i}du_{\kappa',i'}} = \frac{dF_{\kappa,i}}{du} = \frac{F_{\kappa,i}^+ - F_{\kappa,i}^-}{2u} \quad (2.53)$$

where κ is an ion and i is an axis. However, as demonstrated with equation 2.51, the finite displacement method is sufficient for computing only the Γ point if a primitive unit cell is used. For additional wave vectors, it is necessary to model a supercell containing multiple primitive unit cells. Alternatively, phonons may be modelled with the density-functional perturbation (DFPT) method. In DFPT, phonons with arbitrary wave vectors may be modelled with a single unit cell by computing the linear response of the KS orbitals to a displacement of

atoms. All phonon calculations in this work have been performed with the finite displacement method: further description of DFPT is beyond the scope of this thesis.

2.4.2 Vibrational Spectroscopy

Phonons are relevant to structure searching because it is possible to non-invasively measure vibrational frequencies experimentally. IR spectroscopy and Raman spectroscopy are widely implemented and complementary spectroscopic methods that measure a material's response to a fixed-frequency laser. This allows for determination of the frequency of vibrational modes that induce a change in dipole moment (in the case of IR spectroscopy) or a change in electric-dipole polarisability (in the case of Raman spectroscopy).

Photons with a frequency in the IR range can be absorbed by a material to create a phonon of the same energy that is associated with a change in dipole moment. Any phonon that can be created by the absorption of IR photons is said to be IR-active. It is possible to measure all IR-active modes by irradiating a material with a spectrum of IR frequencies and then detecting the frequencies which are reflected or pass through the material. Any material with IR-active modes that is placed between the emitter and detector will absorb some of the IR photons, resulting in fewer photons being detected of certain frequencies. Each absorbed frequency is associated with one or more phonon mode and the amount of absorbed light (the intensity of the signal) gives an indication of the relative number of vibrations which occur at that frequency.

Instead of measuring photon absorption, Raman spectroscopy measures photon scattering. An incident photon may excite an atom to a virtual energy state, which rapidly decays to a lower energy state. This new state may be higher than, lower than, or equal in energy to the initial energy state. If the atom returns to its original energy state, scattered photon shows no change in energy. This is called Rayleigh scattering and it is not particularly diagnostically useful. If the atom falls to an energy state that is *higher* than its starting state

or *lower* than the starting state, the process is called Stokes Raman scattering or Anti-Stokes Raman scattering, respectively, and the change in energy is made up for by the absorption or creation of a phonon. The phonon involved must have an energy equal to the difference in energy between the starting state and the end state, and the associated vibration must induce a change in electric-dipole polarisability. Measuring Raman-active modes is therefore, in principle, as simple as measuring the frequency of scattered light in response to light with a fixed frequency.

Modelling IR and Raman spectra of a structure and then comparing those spectra to experimental spectra can help to determine whether a the structure in the model matches a structure found experimentally.

2.4.3 NMR

Every isotope has a integer or half-integer nuclear spin quantum number, I , which is determined by the nuclear composition. The smallest possible value is 0, but 1/2, 1, and 3/2 are common, with higher values also being possible. Each nucleus for which $I > 0$ can take on $2I + 1$ different magnetic moments and it is this property that gives rise to nuclear magnetic resonance (NMR) [82].

For a nucleus with $I = 1/2$, such as ^1H , there are two possible spin states, $I = 1/2$ and $I = -1/2$, which are degenerate. However, in the presence of a magnetic field, this degeneracy is broken and the energy of the two states is split proportionally to the strength of the magnetic field, \mathbf{B}_0 , the reduced Planck constant, \hbar , and the magnetogyric ratio, γ (which is a constant for each nucleus):

$$\Delta E = \frac{1}{2}\gamma\hbar\mathbf{B}_0 - \left(-\frac{1}{2}\gamma\hbar\mathbf{B}_0\right) = \gamma\hbar\mathbf{B}_0 \quad (2.54)$$

If the nuclei are radiated with photons of energy equal to ΔE , nuclei in the low energy spin state can be excited into the high energy spin state and resonance then occurs between these states. As resonance occurs, the difference in energy is emitted as an EM wave. This energy typically occurs within the radiofrequency range. It is therefore apparent that for a sample of unknown composition, the constituent isotopes can be detected by imposing a magnetic field, radiating with a range of radiofrequencies, and measuring at which frequencies resonance occurs at by detecting EM waves emitted from the sample.

Fortunately, this approach can be extended to identify not only isotopes, but also the chemical environments in an unknown sample. In the solid, liquid, and gas phase, nuclei are surrounded by electrons which influence the strength of the magnetic field experienced by their nuclei. The resultant local magnetic field, B_{loc} , is related to the overall magnetic field, \mathbf{B}_0 , by the magnetic shielding, σ , in the form:

$$\mathbf{B}_{loc} = (1 - \sigma)\mathbf{B}_0 \quad (2.55)$$

but may be expressed more precisely in the form of the Biot-Savart law as:

$$\mathbf{B}_{loc}(\mathbf{r}) = \frac{1}{c} \int d^3r' \mathbf{j}(\mathbf{r}') \times \frac{\mathbf{r} - \mathbf{r}'}{|\mathbf{r} - \mathbf{r}'|^3}, \quad (2.56)$$

where c is the speed of light, \mathbf{r} are electron coordinates, and \mathbf{j} is the electric current in the orbitals induced by the magnetic field. The energy splitting between the spin states is dependent on the shielding constant as,

$$\Delta E = \gamma \hbar (1 - \sigma) B_0, \quad (2.57)$$

and so the resonant frequency is also dependent on the chemical environment. This means that NMR spectroscopy can be used to gain insight into the structure of a molecule/material,

instead of just the isotopes. Because the resonant frequency is dependent on the strength of the magnetic field, NMR spectra are usually reported in terms of chemical shifts so that results are standardised between NMR spectrometers of different strengths. This necessitates the use of a reference compound which is conventionally set with a chemical shift of 0. For each signal, the chemical shift is then calculated, which is defined as:

$$\delta = \frac{\nu - \nu^0}{\nu^0} \times 10^6, \quad (2.58)$$

where ν is the resonant frequencies, and ν_0 is the resonant frequency of the reference compound.

The magnetic field experienced by a nucleus can also be influenced by the spin of neighbouring nuclei, which gives rise to splitting called spin-spin coupling. For a $I = 1/2$ nucleus coupled to another $I = 1/2$ nucleus, there are four possible spin combinations, $+1/2:+1/2$, $+1/2:-1/2$, $-1/2:+1/2$, and $-1/2:-1/2$. In this situation, there are two ways for the spins to be parallel and two ways for the spins to be anti-parallel, giving rise to two peaks with equal intensity. For every additional nucleus included in the spin-spin coupling, the number of peaks increases by one and the ratio of peak intensities change in line with Pascal's triangle. For nuclei with $I > 1/2$, there are more than two possible spin states (there are $2I + 1$ states), and so there is more splitting.

When modelling NMR spectra, neither the resonant frequency, nor the chemical shift are calculated. Instead, only the magnetic shielding is calculated and then an attempt is made to apply an empirical correction to make the result comparable to experimental NMR spectra. This may be done by finding two or more compounds with chemical shifts reported in the literature, computing their magnetic shieldings, and then plotting shift against shielding and finding a line of best fit. The equation to this line of best fit can then be used to convert the magnetic shielding of any other material (within the constraints that the isotopes are the same and the modelling method is the same) to a chemical shift. Therefore an NMR

spectrum can be computed for a proposed structure and compare it to an experimental NMR spectrum as evidence for whether the proposed structure is the same as the one that appears experimentally.

Core electrons make a significant contribution to the shielding and so the pseudopotential and frozen core approximation, as described in Section 2.2.3, do not yield sufficiently accurate predictions. A more accurate description of the wavefunction near the nucleus is needed. One such approach is projector-augmented waves (PAW), which allows for the reconstruction of all-electron wavefunctions from pseudo-wavefunctions. Gauge including project augmented waves (GIPAW), are an extension to this method that has been corrected for wavefunctions in an external magnetic field and introduced a formalism for the electric current in terms of these PAWs.

Chapter 3

LiNiS Electrodes

3.1 Introduction

Conventional LIBs use an LiCoO_2 cathode with a theoretical gravimetric charge capacity of 274 mAh g^{-1} and an LiC_6 anode with a gravimetric charge capacity of 372 mAh g^{-1} [6]. The process of charging/discharging LIBs involves the shuffling of Li^+ (and e^-) between each electrode but a battery cell can deliver no more charge than can be stored by the lowest capacity electrode. Each electrode should therefore be able to accommodate the same number of Li^+ ions, which means that under ideal conditions, 1.36 g of LiCoO_2 are needed for every gram of LiC_6 . Furthermore, in practice, the ratio of LiCoO_2 to LiC_6 will differ from 1.36:1 as the practical charge capacities are always lower than the theoretical capacity, and for LiCoO_2 , the capacity falls by almost half to 140 mAh g^{-1} [83]. Discovering a new cathode material may therefore be an effective path to improve the energy density of LIBs,

Over the past few decades, research has been conducted into several cathodes (such as LiCoO_2 , LiNiO_2 , and LiMnO_2 [83], which are visualised in Figure 3.1) in an attempt to improve the gravimetric (140 mAh g^{-1} [83]) and volumetric capacity that is bottlenecking modern batteries. One of the most promising groups of materials are conversion-type electrodes. One particular material of note is LiNiO (LNO), which has been a topic of

research for as long as LIBs [84–95]. Due to the similar atomic masses of Ni and Co, LNO has a theoretical maximum capacity that is comparable to that of the conventional LiCoO_2 cathode, but fully discharging these materials leads to significant structural degradation and therefore capacity fading. For this reason, it is common to quote a maximum capacity which does not lead to significant degradation. LNO cathodes boast a practical charge capacity of $\sim 200 \text{ mAh g}^{-1}$ [91], 33% higher than that for LiCoO_2 , and have a transition metal framework (Ni) that has a natural abundance 3–4 times greater than Co. The greatest barrier to the practical implementation of LNO batteries is the poor cycling stability (around 20% loss over 100 cycles). This is a consequence of the insertion of Ni(II) into the Li^+ layers, which not only occupies Li^+ sites but also electrostatically hinders Li^+ reinsertion into neighbouring sites. Instability during phase transitions is also known to result in the emission of O from the lattice, which leads to a stoichiometry change, further structural degradation, and charge decay.




Fig. 3.1 Unit cells for LiCoO_2 , LiNiO_2 , and LiMnO_2 , as implemented in LIBs, share the same crystal structure. Li is shown in green, O in red, Co in blue, Ni in grey, and Mn in purple.

Attempts to improve LNO cathode stability by doping with S have been moderately successful [92, 96]. In order to accommodate S atoms, the structure must expand, leading to greater ion diffusivity. The lower electronegativity of S (compared to O) also results in weaker anion- Li^+ interactions, further contributing to the improved diffusivity. The greater inter-layer distance in S-doped materials also results in weakened inter-layer interactions and consequently $\text{Ni}^{3+/4+}$ ions are less likely to migrate into Li^+ vacancies. The improved structural stability of these materials inherently leads to reduced capacity fading. S and O are group 16 elements separated by one period and are known to exhibit similar chemical behaviour. The elements both have p-type valence orbitals and the ability to form 2^- anions, but the ability for sulfur to populate d orbitals [97] also facilitates the formation of 4^+ cations. For this reason it may be possible to synthesise stable LNS structural analogues to known LNO phases, which may have potential as cathode materials. Due to the 100% increase in mass between O and S, one should expect LNS cathodes to have a lower gravimetric capacity than LNO. The enlarged lattice distances which arise from the difference in atomic radii between O and S also comes with the disadvantage of a greater volume and therefore smaller volumetric capacity. Despite this, a number of papers have been published on LNS conversion cathodes and a range of nickel sulfides are known [98–104]. However, there have been no comprehensive studies of the complete LNS phase diagram. AIRSS and the machine learning analogue ‘MLRSS’ has been used to do structure prediction on a broad range of LNS stoichiometries to thoroughly explore the phase diagram and search for LNS materials that can be implemented as cathodes.

3.2 Ternary Phase Diagram

Ternary compounds can take on a range of stoichiometries and geometries. Visualising the full phase space is not straightforward but ternary phase diagrams are a popular and effective way of showing the possible structures made from three different elements. These diagrams

typically take the form of a triangle with a different element on each vertex. Each of these vertices represents structures that only contain that element. The edges of the triangle are effectively binary phase diagrams where the relative position of a point on the edge shows the ratio of atoms of each element in that structure. Each point on the diagram that is not on an edge or vertex represents a structure that contains all three elements, with the proximity to a vertex showing the fraction of atoms in that structure that are of that element. Consequently, a point in the centre of the phase diagram represents a structure with equal quantities of each element. Where multiple structures share the same stoichiometry (polymorphs), the structures share the same position on the phase diagram; though points can be resized or recoloured according to their energy to help differentiate them. Local minima (structures within ≈ 20 meV per atom of the convex hull) should not be neglected entirely as they may be energetically accessible at finite temperatures and multiple polymorphs may be metastable under the same conditions (e.g. C allotropes). Additionally, if the difference in energy between an apparent global minimum and a local minimum is less than the uncertainty from the DFT model, both structures should be considered potential global minima. Ternary phase diagrams do not usually attempt to represent the entire phase space but, rather, are intended to show only global and local minima. Furthermore, given an unlimited number of atoms in a cell, it is possible to define an infinite number of structures with an infinite number of stoichiometries. Some structures are impractical to compute due to the large number of atoms required to define them, and these are therefore necessarily excluded from the phase diagram. Furthermore, where an element makes up a very small fraction of the stoichiometry, the element is effectively a defect and its contribution to the total energy of the cell may be ignored. See section 4.3 for further details on computing formation energies in defects. The number of points on a plot can also be reduced substantially by imposing a maximum energy (usually expressed as a distance from the convex hull), such that particularly high energy structures do not obfuscate structures on the phase diagram that may be stable under relevant

temperatures and pressures. Ternary phase diagrams also commonly include the convex hull, represented as straight lines connecting structures that are thermodynamically stable. Determining which of two polymorphs is more thermodynamically stable is as simple as calculating which polymorph has the lowest total energy. But when comparing two structures with different stoichiometries, their formation energies should be compared. Formation energy is defined as the energy change that occurs when a compound is formed from its constituent elements in their ground states. The formation energy per atom, E_F , of compound, may be expressed as:

$$E_F = E_T - \sum_i X_i E_i, \quad (3.1)$$

where E_T is the total energy per atom of the compound, i is an arbitrary element, E_i is the internal energy per atom of element i in its ground state, and X_i is the fraction of element i in the compound. It is energetically favourable for a material to spontaneously decompose into two or more components if the reaction has a negative change in Gibbs free energy, G ,

$$\Delta G = \Delta E + P\Delta V - T\Delta S, \quad (3.2)$$

where, ΔE is the formation energy and also the change in total energy, P is pressure, V is volume, T is temperature, and S is entropy. Under ambient conditions, pressure and temperature are small, and so the pressure/volume and temperature/entropy contributions are also small (typically on the order of meVs per atom and 10s of μ eVs per atom, respectively). [105] Therefore, it may be approximated that $\Delta G = E_F$, which allows for pressure and temperature to be treated as if they are zero in the DFT model.

Consider a binary system of elements A and B, with an example phase diagram given in Figure 3.2. At one end of the X-axis is pure A and at the other end is B; in-between are $A_{1-x}B_x$ compounds where the distance along the X-axis gives the fraction of A atoms to B atoms. On the Y-axis is the formation energy per atom, which, by definition (see Equation 3.1), is zero for A and B. One can draw a ‘tie line’ between these two points which

acts as a trial for the convex hull. A system of separated A and B atoms has an overall formation energy of zero. But a compound of A and B must have a non-zero formation energy, and if plotted, will appear either above or below the tie line. If it is above, the compound is not stable and it may decompose. If the compound is below the line; it is more stable than its elements. The convex hull may then be redrawn by drawing two new lines connecting this AB compound (with formation energy less than zero) to the elements at either end of the diagram. Each of these convex hull tie lines is a new binary phase diagram with an element at one end, and the compound AB at the other end. Again, for a system divided into some number of AB formula units and some additional number of A (or B) atoms, the formation energy of the box will appear somewhere on the tie line with the ratio between AB and A (or B) determining the position along the X-axis. If a new compound, $A_{0.3}B_{0.7}$, is plotted and it appears above the tie line, it may decompose into some combination of AB and B. But if $A_{0.3}B_{0.7}$ is directly below the tie line, it is more stable than a combination of AB and B, and so the new compound should be incorporated into the convex hull. The process can be repeated for additional compounds and the convex hull is updated until all known compounds have been considered.

The stability of a material is dependent on the formation energy of not only materials with the same stoichiometry but also materials with different stoichiometries, which are represented at different coordinates on the ternary phase diagram. The edges of the ternary phase diagram, which are binary phase diagrams, are of particular interest as, compared to ternary phase diagrams, they are easy to explore thoroughly. The Li-S binary is particularly simple with a single binary compound on the convex hull, Li_2S . This compound is a key feature of a range of LiNiS conversion electrodes. The Ni-S binary is more complex with several binary compounds that are energetically competitive. Of particular interest are NiS_2 , Ni_3S_4 , NiS , and Ni_3S_2 , which as conversion electrodes, have energy densities in the range of $855\text{--}1510\text{ Wh kg}^{-1}$ [98].






Fig. 3.2 Example convex hull for binary compounds with generic formula $A_{1-x}B_x$. Constituent elements, A and B, are shown in blue with formation energy per atom of 0 eV. Compounds on the convex hull are shown in green and compounds above the convex hull are shown in red. The solid black line represents the convex hull, while the dashed black lines represent the hull distance for each compound above the convex hull. In this example, the red compound with formula AB is expected to rearrange to form the green polymorph. The lowest energy compound with formula $A_{0.7}B_{0.3}$ is above the convex hull and therefore the compound is expected to disproportionate into $A_{0.8}B_{0.2}$ and AB.


NiS is one of the most widely studied Ni-S conversion cathode for LiBs [100, 98, 103, 104]. It has a reversible charge capacity of 590 mAh g⁻¹, which exceeds graphite's charge capacity of 372 mAh g⁻¹ by almost double. Additionally, the Li⁺ storage mechanism differs considerably from that of graphite. Intercalation of Li⁺ into graphite results in the formation of a single compound, LiC₆, while storage of Li⁺ into NiS is a two-step conversion process where




occurs at 1.80 V, and the remaining Ni₃S₂ can be converted via



at 1.3 V. NiS₂ is a S-rich material for conversion electrodes with a remarkably high theoretical gravimetric capacity of 712 mAh g⁻¹ [99, 106]. Conversion of NiS₂ into Li₂S and Ni occurs via a similar two-step process but is distinct in that the first step results in the formation of an intermediate with ambiguous stoichiometry. Xu *et al* [106] report that the first step,




occurs at 1.80 V. Which is later followed by



at 1.32 V. The reduction potential reported for NiS₂ is extremely similar to the reduction potential reported for NiS, elsewhere in the literature. If NiS₂ may be reduced to NiS, there must be only a very small difference in reduction potential between the two compounds. The first reduction of NiS₂ is also reported to occur at 1.80 V by Liu *et al* [107] and Jin *et al* [108]

though Jin *et al* ascribe the first reduction to the reaction,



NiS₂ has a number of drawbacks that limits its practical implementation. Of particular note is the material's thermal stability. The material thermally decomposes to NiS and gaseous S₂ through the following reaction:




This decomposition is particularly significant above 550°C, which makes it impractical for use with most molten salt electrolytes. [99] Further to this, NiS₂ has poor conductivity and kinetics, and lithiation/sodiation induces large volume changes, which like other Ni-S materials, can lead to decomposition. Several approaches have been tested to improve stability. Yang *et al* [99] demonstrated that the thermal stability can be significantly improved by increasing the grain size of the crystals. By increasing the grain size from 39 nm to 50 nm, the temperature of decomposition increased by over 100°C. The smallest grain size studied had the greatest discharge performance, but thermal decomposition led to loss in mass and so a grain size that compromises between discharge performance and stability should be used to maximise the gravimetric capacity of the material. Further to this, various substrates have been utilised to alleviate the stability problems that arise from the large volume change. These substrates are frequently C-based, such as reduced graphene oxide (GO) [109–111], 3D graphite oxide (3DGO) [112], C coatings on metal-organic frameworks [113, 106], or C coatings on nanorods [114]. Just like many other materials that struggle with large volume change problems, NiS₂ can also be implemented as nanoparticles to improve structural stability.

NiS₂ is seemingly more widely studied in the context of Na-ion batteries (NIBs). With a Na⁺ charge carrier, it is possible to achieve a charge capacity of 580.6 mAh g⁻¹ at 5 mA g⁻¹ [110], which is lower than the charge capacity reported for the Li system but more than the charge capacity of LiC₆. The conversion process for Na⁺ into NiS₂ is directly analogous to the conversion process for Li⁺ into NiS₂. Li *et al* [110] and Li *et al* [111] report that NiS₂ is firstly reduced to some intermediate Na-Ni-S compound at 1.49 - 1.20 V via




which at 0.82 - 0.89 V can then be further reduced to Ni and Na₂S via



In addition, the redox reactions 3.3 and 3.4 have been reported for both Li⁺ and Na⁺, with the first step for the Na⁺ taking place at 1.39 V [109]. When using a Na⁺ charge carrier, NiS₂ is reported to have a higher reduction potential than NiS. This provides further evidence that when using a Li⁺ charge carrier, the reduction of NiS₂ occurs at a higher potential than the reduction of NiS, but that the difference is small.

For completion, it should also be mentioned that M-S (where M is an alkali metal) conversion electrodes have also been proposed [8, 115, 116]. These electrodes implement S₈ as a starting material and are then converted through to M₂S as:



via one or more polysulfide intermediates (depending on the alkali metal and the electrolyte [115]). In the case of Li₂S, the cathode has a remarkably high theoretical capacity of 1667 mAhg⁻¹, however S₈ suffers from poor cycling performance due to a phase transition

from αS_8 to βS_8 and also decomposition into S_3^- and S_3^{2-} polysulfides which are then lost into the electrolyte. Substrates have been proposed as a means of stabilising the material and thereby improving the cycling performance but this approach has the inherent drawback of reducing the gravimetric charge capacity due to the additional mass of substrate. In order for the cathode to be competitive with LiCoO_2 , the S loading should be no less than 70% of the total mass. Popular host materials of study include graphite and graphene oxide (as mentioned above in the context of NiO), but other materials, such as MnO_2 have also been proposed [8].

3.3 Structure Searching

To efficiently and thoroughly explore the LiNiS phase diagram, a multi-step process has been implemented which is composed of: element swapping, AIRSS, MLRSS, and higher accuracy geometry optimisations.

3.3.1 Element Swapping

As a first step in the structure searching process it is practical to collate structures reported in the literature, and then geometry optimise with the software and settings that will be used for the subsequent calculations. In the case of the Li-Ni-S ternary system, several binary Ni-S structures can be found on online databases but very few are reported for Li-S, Li-Ni, or Li-Ni-S. However, online databases contain considerably more structures which contain elements that are chemically similar to Li, Ni, and S. This can be utilised in a process called element swapping. A search was performed across the Materials Project [117] and Open Quantum Materials Database [118] for ternary compounds that contained one alkali metal, one transition metal, and one chalcogen, in any ratio. All structures were downloaded, and swapped alkali metals for Li, transition metals for Ni, and chalcogens for S. Geometry

optimisations were then performed on all structures, and these form the basis of the ternary phase diagram.

Elements from the same group of the periodic table are chemically similar (compared to elements from outside their group) and it is expected that regions of phase space that are energetically favourable for one combination of elements are also favourable for a similar set of elements. This allows for leveraging of structure searching performed by other researchers on a broad set of elements as a qualitative approximation for structure searching within another particular system. However, any element substitution will have a non-negligible impact on the formation energy of a structure. The chemical trends within a group are indeed *trends*, which is to say that elements neighbouring on the periodic table have similar but not identical chemical behaviours. Generally speaking, the greater the difference in period, the less similar two elements within a group will be. In most cases, this effect largely arises from the difference in the ionic radius that inherently arises from a greater number of electrons. For some groups, there are considerable anomalies in these trends. In the case of the chalcogens, O typically only occupies a 2^- oxidation state, while S, Se, and Te have all also been observed to take the 2^+ , 4^+ , and 6^+ oxidation states. Furthermore, O, S, and Se are usually considered non-metals, while Te is sometimes called a metalloid, and Po a post-transition metal. The range of oxidation states taken by transition metals is also quite broad and depends on their group (of which there are 10 for transition metals), but they are more similar to each other than they are to non-transition metals. Group 1 elements have their own irregularities. H is notoriously anomalous and although usually placed in group 1, it is never considered an alkali metal (and was not included in the element swapping). Furthermore, the alkali metals notably behave differently when intercalated into graphite, with Li saturating at a ratio of 6:1, K at a ratio of 8:1, and Na at a considerably higher ratio for reasons that aren't fully understood. For all groups, additional factors also come into play for heavier elements, such as radioactivity and special relativity. It is for these reasons

that no attempt is made to make a quantitative estimate of formation energy for element swapped structures. However, as geometry optimisations are performed on these structures, the formation energy for these structures can be calculated as accurately as the Li-Ni-S structures sourced directly from a database (although, in practice, Li-Ni-S structures are also subject to geometry optimisations as it is unlikely that the database structures were optimised with the same software and settings as used in this research).

Like many other structure searching approaches, this method is likely to yield several duplicate structures. Automatically determining whether two cells contain identical/equivalent structures is challenging and so it is expected that databases contain a small number of duplicate structures. But additionally, some structures are likely to appear in a database multiple times across different stoichiometries and when element swapped to Li, Ni, and S, become the same cell (but with some volume differences).

3.3.2 AIRSS and MLRSS

The data set of Li-Ni-S compounds was next expanded by performing random structure searching, firstly *ab initio* random structure searching (AIRSS) and later machine learning random structure searching (MLRSS). No attempt is made to explore all possible stoichiometries but instead the stoichiometries that are explored are chosen selectively. Limits are imposed by the large cost of performing geometry optimisations on cells containing a large number of atoms, as the minimum number of atoms needed to define the cell is dependent on the stoichiometry. Therefore, cells which contain an extremely large or extremely small ratio of one particular element become prohibitively expensive to compute. Additionally, the preliminary phase diagram, made from element swapping and also structures that are known for similar systems, can be used to identify regions that are unlikely to yield low energy structures. In particular, structures that are mostly Li are avoided. In addition to exploring specific stoichiometries, particular focus is given to exploring the tie lines connecting Li,

and either NiS_2 , Ni_3S_4 , NiS , or Ni_3S_2 as these lines represent the range of stoichiometries achieved by addition of Li to each of these Ni-S binaries.

The first step of random structure searching is to generate a random cell with a set of constraints imposed by the user. One such constraint is the number of atoms of each element that are to be included in the cell. In the case of specific stoichiometries, the ratio of each element is fixed and an integer number of formula units between 1 and 4 are randomly generated. Searching a range of formula units is important as a structure can only be found if there is a common factor between the number of formula units in the primitive cell and the number of formula units in the search. In order to explore a tie line from a Ni-S binary, the ratio of Ni to S is fixed, a random integer number between 1 and 6 of Ni-S formula units are added to the cell, and then between 1 and 4 Li atoms are added to the cell. This more complex approach allows me to generate a range of stoichiometries and a range of formula units simultaneously. Several constraints were then set up, including the crystallographic space group (which is randomly selected), minimum distances between atoms (slightly less than typical bond distances), and the volume (randomly selected between 15 \AA^3 and 20 \AA^3 per atom). Cell parameters and atomic coordinates are then generated randomly with the atoms provided and so as to satisfy the imposed constraints.

Geometry optimisations were then performed on the generated structures in order to minimise the energy, forces, and stress. These geometry optimisations are performed with the same parameters used for the element swapping geometry optimisation. In order to minimise the computational cost of these optimisations, and hence increase speed at which the phase space is explored, accuracy DFT parameters are chosen that are accuracy and cost efficient. The specific parameters are detailed in Table 9.1. The formation energy of the structure can then be calculated from the total energy of the optimised structures, and it can be added to the phase diagram.

The time required for AIRSS to generate structures is largely dominated by the geometry optimisation step, which is itself dominated by performing force evaluations with DFT. In an attempt to accelerate the structure searching process, DFT is replaced with a machine-learning potential (MLP) for computing energies/forces/stresses in a process termed ‘Machine Learning Random Structure Searching’ (MLRSS). MLRSS was not possible as a first step as the first step of the structure searching process as no Li-Ni-S MLP was available. Therefore a new potential was trained. Performing element swapping and AIRSS allowed me to simultaneously populate a phase diagram and also generate a database of Li-Ni-S structures that could then be used to train the ANN. The initial approach to MLRSS was an iterative one, where a small number of structures was used to train a crude MLP, use that to perform geometry optimisations, and then use DFT to compute energies/forces/stresses for the MLP-optimised structure. This data could then be used to expand the training set that could be used to retrain the MLP in order to improve the accuracy of its predictions. The initial implementation of this approach was flawed as the poor force predictions from the initial MLP resulted in structures that contained atoms with exceptionally short bond distances. In some situations, where nuclei appeared within the radius of another nucleus’ pseudopotential, CASTEP returned inaccurate energies, forces, and stresses. This further exacerbated the poor force predictions when this inaccurate data was used to train the second generation MLP. After discovering this problem, geometry filters were implemented to the training set that removed all structures with inaccurate energy/force/stress values. Third generation and later MLPs therefore did not suffer from the critical flaw introduced to the second generation MLP. In addition to the geometry filters, energy, force, and stress filters were also implemented to better focus the ANN and improve the quality of the geometry optimisations for most of the explored phase space. By the time the fourth generation MLP was trained, the MLPs predicted forces with sufficient accuracy, but not stresses. Subsequent generations of MLPs were therefore trained on stresses (in addition to forces) but no stresses were available for the

MLP-optimised structures as stresses are not computed by default for singlepoint calculations in CASTEP 19. Consequently, the fourth generation MLP was trained entirely on element swapping and AIRSS data. At the point of the fourth generation, the MLP suffered from a problem of overfitting, which is to say that the MLP was able to learn the properties of structures in the training set but poorly reproduced the properties of structures outside of this set. This is largely attributed to the database, which had a bias towards particular stoichiometries that had been focused during AIRSS (the NiS_2 -Li tie line in particular) and also a lack of structures with large forces. To address this problem, a new database of structures was generated. This was initiated by generating a small number of AIRSS structures but performing a DFT singlepoint calculation on the starting structure instead of performing a geometry optimisation. This DFT data was then used to assemble the first generation of the new database and trained a PaiNN MLP. The first generation MLP was then used for a small number of MLRSS runs, and further DFT singlepoint calculations were performed on the “optimised” structures. The force errors in the first generation MLP were large and so the expectation was that the forces in many of the structures, as modelled by DFT, would be large. By adding this new DFT data to the database, the subsequent generations of MLPs can learn that the forces are large for these structures and therefore the MLP-driven geometry optimisations become more reliable. The process is repeated until the force and stress error is sufficiently small that MLRSS regularly produces structures that when modelled with DFT, return small forces and stresses. A summary of the MLRSS process is shown in Figure 3.3. The learning curve shown in Figure 3.4 shows how the force error in each validation set reduces with respect to the number of structures in the training set. The latest generation MLP may then be used for further MLRSS until the ternary phase diagram is sufficiently populated.

Every structure generated by MLRSS is subjected to a DFT singlepoint calculation in order to find the formation energy. For any structures that are close to the hull, an additional

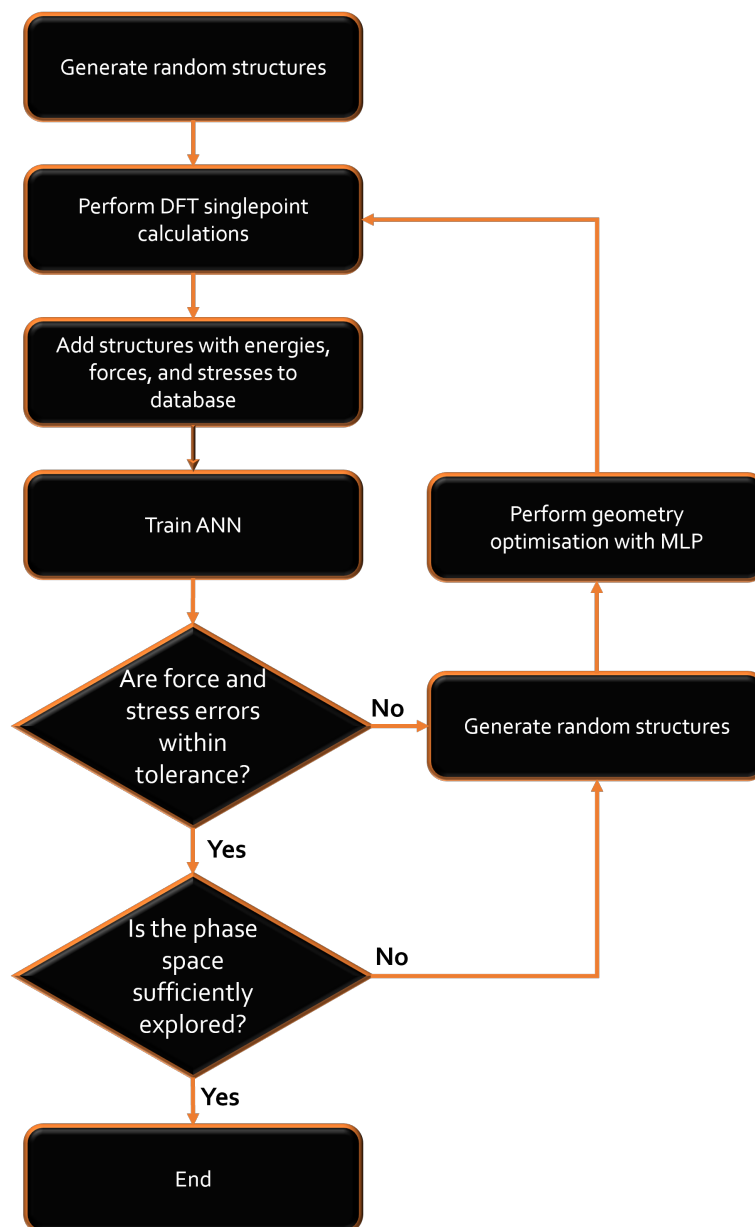


Fig. 3.3 Flow chart to show the iterative process by which MLRSS simultaneously explores a phase space and improves force and stress errors of the MLP. All geometry optimisations are performed with an MLP. If force and stress errors are within an acceptable tolerance, further training of an ANN is not necessary, but depending on the available compute time, may be an effective means to further improve the performance of MLRSS.

DFT geometry optimisation was performed. See section 3.3.4 for further details. The approach used to acquire the latest MLPs does not differ significantly from the original approach employed in this research but a few, seemingly small, modifications were made to the method which are believed to be responsible for the improved performance. Firstly, geometry filters were implemented before any DFT calculations were performed. This acted to ensure that inaccurate DFT data did not enter the training set and prevented compute time being spent on DFT calculations that could not be used. Secondly, the number of atoms that could be included in an AIRSS cell was reduced. Early AIRSS runs generated cells containing a few dozen atoms, depending on the stoichiometry, but by constraining the cell to between 3 and 12 atoms, the time required for any DFT calculations is reduced, allowing for the database to be expanded faster and also introducing more variety to the database (as fewer structures contain multiple similar/equivalent atoms). The search space was also expanded to include any structure with stoichiometry $\text{Li}_x\text{Ni}_y\text{S}_z$, where x , y , and z may take values 1-4; this further acted to add more diversity to the training set. Previously, a maximum force filter was applied to the training sets such that structures with forces greater than 4 eV / Å but this tolerance was raised to 50 eV / Å which allows the MLP to learn a more accurate description of the phase space navigated during earlier steps of the geometry optimisations, and further expanded variety in the training set. Lastly, there was a shift from large increases in database size between generations to small increases in database size between generations. This is believed to allow the MLPs to more dynamically react to errors, rather than making the same “mistake” hundreds of times.

3.3.3 Element Swapping with AIRSS

Element swapping and AIRSS was also used in combination as part of a limited approach for expanding the explored phase space. Under my supervision, an undergraduate student, Mr. Wei Huang, conducted a summer research project as part of the Morris Group in which

he used AIRSS to explore the Ni-S binary phase diagram and a small number of important stoichiometries of key interest on the M-Ni-S ternary phase diagram, where M is Li, Na, or K. The Na and K atoms were element swapped to Li, and the all the structures discovered by Mr. Huang underwent further geometry optimisations with the parameters given in Table 9.1. In addition, AIRSS was recently applied to explore the Li-Ni-O ternary phase diagram by Cen *et al* [95]. Several Li-Ni-O ternary compounds were discovered on the convex hull that had not been reported on the Materials Project [117]. Element swapping was performed on each of these structures and added them to this project's dataset.

3.3.4 Higher Accuracy Geometry Optimisations

As mentioned in section 3.3.2, the initial DFT calculations performed as part of the structure searching, were chosen as a compromise between accuracy and cost. The accuracy of the phase diagram is next improved by performing further DFT geometry optimisations on the structures closest to the hull but with more accurate DFT parameters. Notably, spin polarisation is enabled, the k-point spacing is reduced from 0.05 \AA^{-1} to 0.04 \AA^{-1} , and a more accurate pseudopotential is used. The important parameters are summarised in Table 9.3. Further optimisations are performed only on the structures within 100 meV / atom of the convex hull, as the change in formation energy is rarely greater than 100 meV / atom. Figure 3.5 shows the difference in formation energy after geometry optimising over 2500 structures within 100 meV / atom of the convex hull. Fewer than 10 structures changed in formation energy by more than 100 meV / atom, with the majority of structures changing by less than 20 meV / atom.

Although each SCF cycle is more expensive for the higher accuracy geometry optimisations, the structures are already close to their local minima and so fewer geometry optimisation steps are needed. This additional step ensures that all structures on the phase diagram have been geometry optimised to a DFT level of accuracy, whereas the hull without

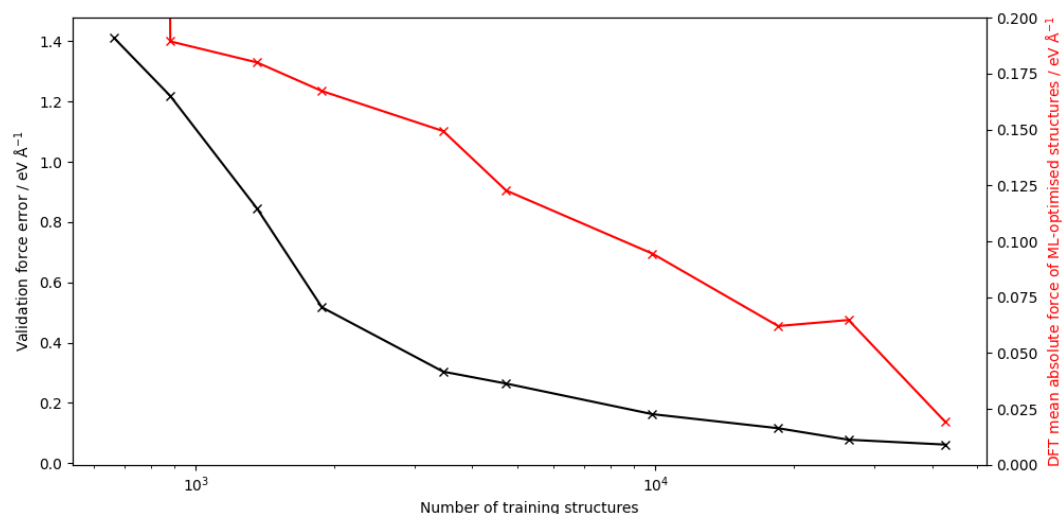


Fig. 3.4 Reduction in force errors in validation set for subsequent generations of MLPs. X-axis shows number of structures in the training set. Left Y-axis with black curve shows the error in the validation set for each MLP. Right Y-axis with red curve shows the mean absolute force, as modelled by DFT, of structures geometry optimised with each generation of MLP. Additional structures in the training set are sourced by performing MLRSS with the latest generation of MLP and performing a DFT singlepoint calculation on the optimised structure.

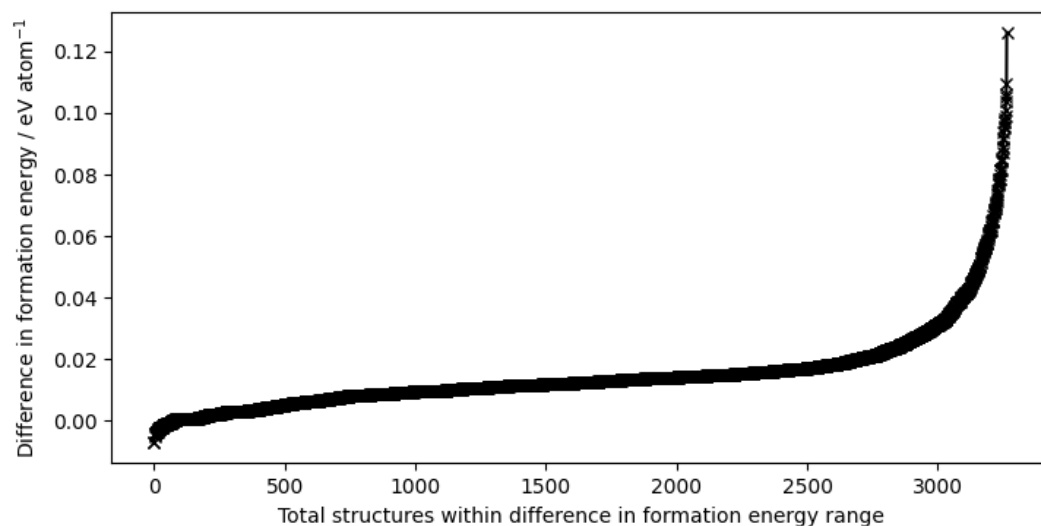


Fig. 3.5 The difference in formation energy after higher accuracy geometry optimisations of structures within 100 meV / atom of the convex hull. The X-axis shows the number of structures with a difference in formation energy per atom less than the value on the Y-axis.

this step included some structures that had only been optimised with the less accurate MLP forces.

3.3.5 Ni-S Binary Convex Hull

The Ni-S binary convex hull was investigated first as this binary system effectively represents the possible uncharged state of a potential Li-Ni-S electrode. Unlike the Ni-Li hull and the S-Li hull, multiple binaries are known to be stable and each of these binaries have multiple polymorphs within a few meV. It therefore seems more likely that there may be undiscovered but energetically stable polymorphs, so considerably more time was spent exploring this binary than the others (much of which was performed by Mr. Huang (see Section 3.3.3). This hull, shown in Figure 3.6, makes it clear that Ni_3S_2 , NiS, Ni_3S_4 , and NiS_2 all have multiple energetically competitive polymorphs. While Ni_3S_2 and Ni_3S_4 have previously been reported on online databases [117, 118], new structures are reported for NiS and NiS_2 that have been discovered through AIRSS. Each of these minima are shown in Figure 3.7. Ni_3S_2 and NiS_2 are, on visual inspection, interesting candidates for a battery material as they are likely to have good Li^+ kinetics due to the channels in the prior and the layered structure of the latter.

The minima reported here are for a system at 0 K and 0 GPa but any practical LiNiS cathode would be expected to be implemented under atmospheric conditions. At finite temperature, the $T\Delta S$ term of the Gibbs free energy (see equation 3.2) becomes non-zero. In the absence of phonon modelling, ΔS has not been quantified, however, one can estimate an energy window within which polymorphs may become stable under a given temperature from the $k_B T$ energy. At 298 K, this energy is approximately 26 meV. For the four Ni-S binaries modelled to be stable at 0 K, 4 additional polymorphs were identified for NiS (visualised in Figure 3.9), 2 additional polymorphs for Ni_3S_2 (visualised in Figure 3.8), and 0 for either NiS_2 or Ni_3S_4 .




Fig. 3.6 Convex hull for binary Ni-S system, showing Ni_3S_2 , NiS , Ni_3S_4 , and NiS_2 on the convex hull.




Fig. 3.7 The lowest energy structures discovered for Ni_3S_2 (a), NiS (b), Ni_3S_4 (c), and NiS_2 (d). The polymorphs of NiS and NiS_2 have not reported on the MP, OQMD, or ICSD databases. S is shown in yellow and Ni in grey.





Fig. 3.8 Visualisation of Ni_3S_2 polymorphs that are above the convex hull and within 26 meV / atom of that hull. Distance from the hull is shown above the polymorph in meV / atom. S shown in yellow and Ni in grey.

3.3.6 Li-Ni-S Ternary Convex Hull

Despite generating over 150,000 structures via AIRSS and MLRSS, no ternary compounds were identified on the convex hull. The ternary compound closest to the hull, which had stoichiometry LiNi_2S_2 (see visualisation in Figure 3.10), was 43 meV / atom away and is therefore unlikely to be stable at ambient temperature. Further, LiNi_2S_2 is on the NiS-Li tie-line and if it were stable, it is expected that it would be observed experimentally during lithiation of NiS cathodes [100, 98, 103, 104].

Considering the complete ternary phase diagram, Figure 3.11, there are only eight stable phases, three are the ground states of the constituent elements, Li, Ni, and S, four are on the Ni-S binary hull, Ni_3S_2 , NiS, Ni_3S_4 , and NiS_2 , and the remaining phase is the Li-S binary, Li_2S . Five of these phase appear suitable for lithiation. The first is a simple S cathode, which may accommodate Li^+ via the reaction given in equation 3.11, in which the Li_2S has a gravimetric charge capacity of 1167 mAh g^{-1} . The alternatives are Ni_3S_2 , NiS, Ni_3S_4 , and NiS_2 , which, as discussed in Section 3, can be utilised as conversion cathodes. These materials have theoretical gravimetric charge capacities of 400, 512, 595, and 712 mAh g^{-1} , respectively. In practice, these four binary compounds may be viewed as a single conversion cathode as lithiation of higher S-content Ni-S compounds leads to disproportion to lower S-content Ni-S compounds, and eventually Ni, via the multiple step reaction:



A NiS_2 conversion cathode of this type is modelled to show a reduction potential of 2 - 1.5 V, as shown in Figure 3.12. The reduction potential of 1.8 V for step 3 of the above equation matches very strongly with the widely experimentally reported reduction potential of 1.8 V.




Fig. 3.9 Visualisation of NiS polymorphs that are above the convex hull and within 26 meV / atom of that hull. Distance from the hull is shown above the polymorph in meV / atom. S shown in yellow and Ni in grey.




Fig. 3.10 Visualisation of LiNi₂S₂ (left): the ternary compound closest to the convex hull with a distance of 43 meV / atom. Visualisation of LiNiS₂ (right): the Li_xNiS₂ compound closest to the convex hull with a distance of 56 meV / atom. Li shown in green, S in yellow and Ni in grey.

The reduction potential for step 4 is slightly overestimated with a value of 1.5 V against the experimentally reported 1.3 V [98, 100, 101, 103, 104]. For NiS_2 , a reduction potential of 2.0 - 2.2 V is modelled, which differs by 0.2-0.4 V from the value of 1.8 V reported by Xu *et al* and Jin *et al*. An error of 0.2-0.3 V can typically be expected for a DFT model [119–122], and given this consideration, these calculations are generally consistent with experiment.




Fig. 3.11 Convex hull for ternary Li-Ni-S system shows no ternary compounds appear within 40 meV / atom of the convex hull. All stable compounds appear on the vertices or on the Ni-S and Li-S binary convex hulls. The three vertices show pure Ni, S, and Li in their ground states. The Ni-S axis features Ni_3S_2 , NiS, Ni_3S_4 , and NiS_2 on the convex hull. The S-Li axis features only Li_2S on the convex hull.

3.3.7 NiS_2 -Li Tie Line

The model Li-Ni-S ternary phase diagram, Figure 3.11, suggests that reduction of NiS_2 should proceed via the four-step reduction in equation 3.12. A cyclic voltammogram of a NiS_2 cathode is therefore expected to show four reduction peaks, which, given the voltage curve, Figure 3.12, should appear at 2.00 V, 1.88 V, 1.80 V, and 1.53 V. However, this

contrasts against the two-step reduction observed by Xu *et al* at 1.80 V and 1.32 V which they claim are related to reduction of NiS_2 to Li_xNiS_2 , and reduction of Li_xNiS_2 to Ni and Li_2S , respectively [106]. A reduction potential of 1.80 V is also consistent with the experimentally observed and modelled reduction potential of step 3 in equation 3.12. If NiS_2 can be reduced to Li_xNiS_2 , a Li_xNiS_2 compound is expected to be found on the Li-Ni-S convex hull, but no such compound is identified. Figure 3.13 shows a plot of all Li_xNiS_2 compounds discovered, with the ratio of Li to NiS_2 on the X-axis and the hull distance on the Y-axis. The structure closest to the convex hull had a hull distance of 56 meV / atom and the chemical formula LiNiS_2 (see Figure 3.10 for visualisation). Compounds with 15 different stoichiometries were identified within 200 meV / atom of the convex hull, while only 1 was identified for $\text{Li}_x\text{Ni}_3\text{S}_4$, 3 for Li_xNiS , and 0 for $\text{Li}_x\text{Ni}_3\text{S}_2$, suggesting that Ni: S_2 is a “magic ratio”, leading to unusually stable Li-Ni-S ternary compounds. The lowest energy polymorph of each stoichiometry had a mean density of 3.37 g / cm^3 , with a maximum and minimum density of 3.83 and 2.32 g / cm^3 , respectively. AIRSS/MLRSS are not comprehensive structure searching approaches and it remains possible that a Li_xNiS_2 structure exists on the Li-Ni-S convex hull but was undiscovered. The structure reported by Xu *et al* and the structure reported by the Materials Project as the groundstate for Li_xNiS_2 is a cubic pyrite-type compound, while the ground state identified in this work is a monoclinic layered compound, visualised in Figure 3.7. The cubic pyrite-type structure appeared 59 meV / atom above the convex hull, 3 meV / atom greater than the hull distance of the closest ternary compound. The large hull distance of cubic pyrite Li_xNiS_2 may explain the thermal instability of the compound, and if this polymorph of NiS_2 is metastable, the ternary Li_xNiS_2 compounds identified at similar hull distances may also be metastable.

Discovering the structure may necessitate larger numbers of atoms in the AIRSS/MLRSS cells and a broader range of stoichiometries.

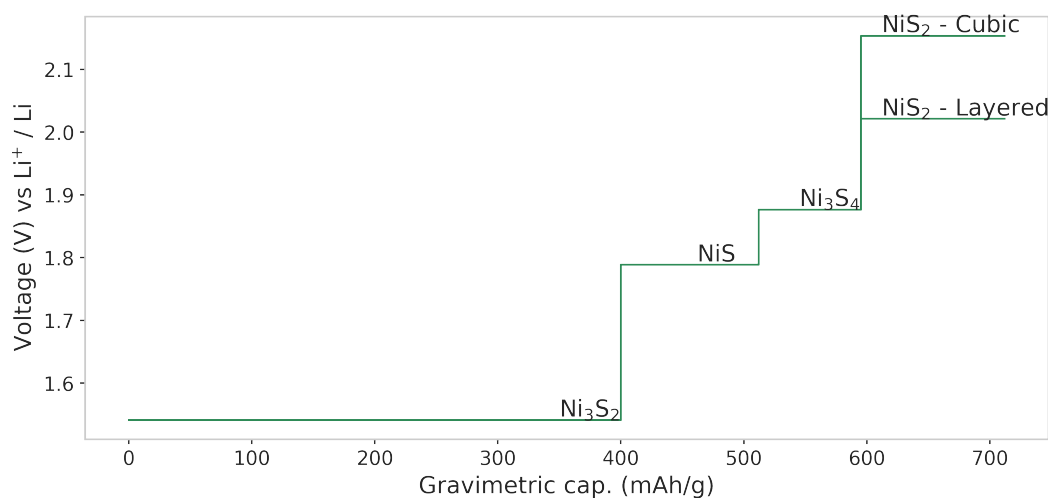


Fig. 3.12 Voltage curve across lithiation of NiS_2 and conversion into Ni_3S_4 , NiS , Ni_3S_2 , and Ni , alongside Li_2S . See equation 3.12 for the stepwise conversion reaction. Potential matches strongly with the experimentally reported reduction potentials of ~ 1.80 V for NiS and ~ 1.30 V for Ni_3S_2 [101, 104, 103, 100, 98]. A reduction potential for NiS_2 is reported for the lowest energy polymorph found in this work, NiS_2 - Layered, at 2.02 V and the NiS_2 polymorph reported in the literature, NiS_2 - Cubic, at 2.16 V. These values correlate well with the 1.9 - 2.5 V potential reported for unlithiated NiS_2 cells [106–108].

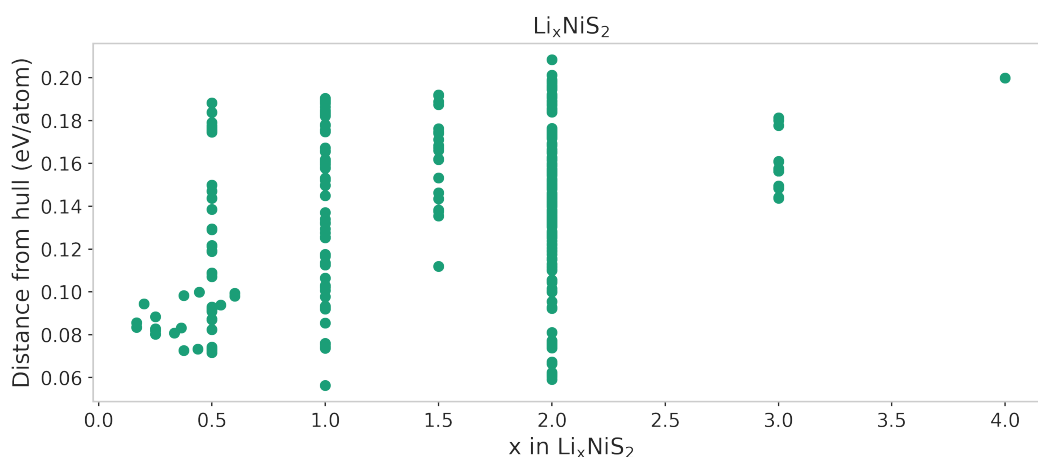


Fig. 3.13 Distance from Li-Ni-S convex hull for each discovered structure with general formula Li_xNiS_2 . X-axis shows ratio of Li to NiS_2 in compound, while Y-axis shows distance from the Li-Ni-S convex hull. The lowest energy structure has stoichiometry LiNiS_2 but is 56 meV / atom above the convex hull; it is therefore energetically unstable at room temperature.

3.4 Conclusion

Element swapping from online structure databases, local structure databases, and a Li-Ni-O structure searching publication [95], and AIRSS/MLRSS were performed to search Li-Ni-S phase space for stable Li-Ni-S ternary compounds with a high gravimetric charge capacity such that it may be applied as a cathode in Li-ion batteries. The ternary phase space was explored thoroughly but no stable ternary compounds were identified. A Li-Ni-S phase diagram has been plotted that, to the author's knowledge, is the most complete in the literature, and a computed a voltage curve for Li-NiS₂ and NiS based conversion cathodes have also been plotted. While the model for Li-NiS based conversion cathodes is consistent with the experimentally observed electrochemical behaviour of Li-NiS conversion cathodes, the model shows inconsistency with the experimentally observed behaviour of Li-NiS₂ based cathodes. The polymorph of NiS₂ observed experimentally has a cubic pyrite-type structure but a layered polymorph of NiS₂ has been identified that is 58 meV / atom lower in formation energy. Xu *et al* [106] report an intermediate, Li_xNiS₂, during lithiation and while no ternary compounds appear on the convex hull, Li_xNiS₂ compounds were discovered with a similar distance from the convex hull to the cubic polymorph of NiS₂. Xu *et al* may have observed a metastable Li_xNiS₂ compound, but the possibility that a stable Li_xNiS₂ compound exists in Li-Ni-S phase space and has not yet been discovered cannot be eliminated. Additionally, cubic NiS₂ thermally decomposes into S₂ and NiS, which suggests that cubic NiS₂ and Ni₃S₄ (since it is between NiS₂ and NiS on the Ni-S binary phase diagram) are not on the hull. A convex hull is reported for Li-Ni-S at 0 K and 0 GPa, but it is possible that under ambient conditions, Ni₃S₄ move away from the convex hull as the PV and TΔS components of Gibbs free energy (equation 3.2) become non-zero.

Further structures searches may lead to a more complete phase diagram and convex hull. The phase diagram may now be populated at a faster rate than it could have been previously as this iterative approach to MLRSS has lead to progressively more accurate MLPs that can

perform geometry optimisations three orders of magnitude faster than DFT. Furthermore, training each of the MLPs took fewer than 6 hours on GPUs, which is negligible compared to the compute time required to geometry optimise enough structures to populate a ternary phase diagram. Geometry optimisations will be less accurate than DFT geometry optimisations, but if higher accuracy geometry optimisations would otherwise be performed with AIRSS, this slight loss in accuracy is of little impact as it will not be carried through to the structures used to construct the convex hull. Given the large gains in computational efficiency, and the negligible disadvantages of MLRSS, it appears to be a powerful approach for any future large-scale structure searching tasks.

Chapter 4

Monolayer hBN Defects

4.1 Introduction

Hexagonal boron nitride (hBN) is a layered hexagonal polymorph of boron nitride, with a similar structure to graphite. Just as 2D atomically-thin monolayers of graphite (graphene) can be synthesised through low-pressure chemical vapour deposition, and exfoliation, hBN monolayers can be synthesised by the same methods [123]. In 2015, hBN monolayers were demonstrated to behave as a single-photon emitter (SPE) at room temperature [124]. Several SPEs are known, including single atoms [125], single molecules [126], quantum dots [127], and carbon nanotubes [128], but at present, hBN is the only 2D material known to act as an SPE above cryogenic temperatures [124, 123].

SPEs are materials that can only emit a single photon at a time, which is to say that a minimum amount of time must elapse between each photon emission. They are therefore a means of generating a single photon of a fixed frequency. In contrast, classical light sources, whether coherent (like lasers) or incoherent (like lamps) release multiple photons together. For a material to be considered an SPE, it must be capable of emitting a photon such that a given detector, which can determine the number of incident photons, should measure only one photon for each incident state. Furthermore, if the photon is directed towards a beam

splitter and a detector is placed at the end of each possible path, a photon should only be measured on one detector. Photons passing through a beam splitting have a 50% chance of being transmitted through the splitter and a 50% chance of being reflected. For a continuous light source, the light beam is therefore split evenly in intensity between the two paths, but for a single photon the “beam” may only take one path [129].

SPEs have a range of possible applications in quantum technology, such as quantum computing (where the polarisation of a single photon may be used as a qubit) and quantum cryptography (where quantum properties of photons may be used as cryptographic keys) [130–136]. They are also important for calibrating single photon detectors, which have their own applications, which also includes quantum computing and quantum cryptography [129].

The source of the single photon emission in monolayer hBN is not known exactly but it has been attributed to point defects. Consequently, an important step towards engineering SPE instruments is identifying the defect(s) responsible so that they can be synthesised for implementation. A range of defects have been studied, including atomic substitutions with heteroatoms, B or N atoms in B or N sites (conventionally expressed as {element}_B or {element}_N, respectively), or vacancies in B or N sites (conventionally expressed as V_B or V_N, respectively), and typically in the range of one to three defect sites [123, 137–144]. Prior research has largely focused on some combination of B, N, and C substitutional defects and vacancies, as these are generally considered to be the most likely sources [123, 139–144]. This is due to the frequent exposure of hBN to C during synthesis and the experimentally observed correlation between photoluminescence intensity and the exposure of the material to C. The same correlation is not observed for other common defects such as Si and O [137]. Models of photoemission from C defects have also previously been shown to correlate well with experiment [137]. However, the literature covers only a narrow range of defects with most publications discussing fewer than ten unique defects, which have likely been individually selected. A targeted approach such as this leaves a large fraction of phase space

unexplored, potentially allowing for defects with useful optical properties to go overlooked. Therefore, an attempt is made to expand the number of considered defects through high-throughput defect searching. Formation energies are then computed as a measure of stability. This approach yields hundreds of unique defects, including defects which extend beyond the plane of the monolayer (extrinsic defects) and various polymorphs of many intrinsic defects, defects that occur entirely within the hBN plane.

This research was conducted in collaboration with Siyu Chen and Prof. Bartomeu Montserrat from the Materials Theory Group at the University of Cambridge. Mr. Chen and Prof. Bartomeu Montserrat will be utilising Yambo [145] modelling excitons and the optical absorption and emission spectra for low energy defects with DFT+GW and the Bethe-Salpeter equation (BSE) [146, 147]. Excitons, and optical absorption/emission spectra are dependent on the excitation of an electron from the valence band to the conduction bands. Consequently, the accuracy of an exciton model is dependent on the accuracy of the band gap in the model, and for DFT+LDA and DFT+GGA the band gap is underestimated. The accuracy of the band structure can be improved through the introduction of the GW approximation via the use of the Green's function, G , and the screened Coloumb interaction, W . Further, the BSE may be solved in terms of Green's functions in order to compute the density response to the excitation of an electron and the subsequent generation of an exciton. The density response impacts the Coloumb screening experienced by the exciton and hence its binding energy which determines the probability of the creation or destruction of an exciton: which relates directly to the intensity of a peak in an optical absorption or emission spectrum. The optical absorption and emission spectra produced by Mr. Chen and Prof will be compared to experimental measurements of single photon emission to identify which of the proposed defects are likely candidates for single photon emission.

4.2 Defect Searching

An AIRSS approach has been taken to searching for low energy defects in hBN. AIRSS is a highly versatile structure searching algorithm with paramaterisation pertaining only to the generation of the starting structures. AIRSS allows for precise control of the types of structures that are being explored and is therefore well suited to defect searching [36–38] as only a small area of a cell and only a small number of atoms deviate from the pristine cell. Defects, by their nature, are aperiodic, so a sufficiently large supercell must be used in order to minimise self-interactions across the periodic boundaries of the cells. To this end, all structure searching was performed in a 5×5 supercell with a 12.5 \AA width and a 20 \AA vacuum region in the Z-axis. Although the effect of self-interaction is further minimised with larger supercells, an increase in the size of the cell also increases the computational cost of DFT calculations and so 5×5 has been selected as a compromise as cells of this size are sufficient to mitigate interactions between defect sites in periodic images. A spherical defect search area is defined within each cell with radius 2.5 \AA , within which all defect searching is contained. The search area is centred over a hBN hexagon to ensure that there are a whole number of formula units both inside and outside, which helps reduce bias in the stoichiometry of the searched structures. The search area contains three hBN formula units, which provides sufficient flexibility for generating a range of defects, and the search area is 7.5 \AA from the nearest point of the search area in the repeating cell to minimise defect self-interaction. The $5 \times 5 \times 1$ pristine hBN cell and the search area can be seen in Figure 4.1. The cell parameters were fixed during structure generation and geometry optimisations to model the constraints imposed on the cell by the bulk structure. This approach has the disadvantage of straining the cell, which increases the energy of the defect. Mitigating the strain is possible through the use of larger cells, but as mentioned previously, this would result in more computationally demanding geometry optimisation. All intrinsic defects and the lowest energy defect of each unique stoichiometry and number of atoms were later embedded in a larger $15 \times 15 \times 1$

and geometry optimised again. The difference in formation energy (defined in section 4.3) for a pseudocell in a $5 \times 5 \times 1$ supercell and larger supercells was approximately 136 meV, significantly less than the total energy range of 40,000 meV, and an order of magnitude smaller than the typical energy difference between low energy competing defects. Therefore, constraining structure searching to $5 \times 5 \times 1$ supercells is an appropriate approximation for the initial geometry optimisations. Further details concerning defects in $15 \times 15 \times 1$ cells can be found in Section 4.4




Fig. 4.1 Pristine $5 \times 5 \times 1$ hBN cell used for generating defects with AIRSS. The defect searching area is marked in a black circle.

All structure searching was performed with AIRSS as implemented in an in-house Python 2 package. The defect searching included randomly removing B and N atoms and randomly adding B, N, and C atoms in such a way the net change was never greater than three atoms. This constraint was imposed to reduce the size of the explored phase space, to minimise the complexity of the geometry optimisations, and because larger defects are generally higher energy and hence less likely to occur naturally. Searching was also split into two parts; one in which an atom could be inserted above/below the plane of the monolayer, and one in

which changes to the cell were constrained to the hBN plane. The latter constrains searches to a smaller phase space but the generated cells are, on average, much faster to geometry optimise and defects with atoms outside of the plane of the monolayer are generally higher in energy. Defect searching continued until that phase space had been sufficiently explored, as indicated by consistently generating defects which had had been discovered earlier in the AIRSS process, with simple defects appearing over 50 times. A very limited number of defects were constructed manually to ensure that the defects that had already been reported in the literature were considered. However, it was also verified that these structures were discovered independently by AIRSS. Figure 4.2 summarises the energy distribution of all structures discovered through AIRSS. Instances where a defect was found multiple times appear on the plot as a vertical stacks of crosses. These are most prominent for simple defects and non-defected hBN (at 0 eV), which are found more frequently. This is represented more clearly in Figure 4.3, which shows the number of structures generated of each stoichiometry. Due to the constraint placed on the structure searching that each defect should deviate from the pristine by no more than three atoms, the explored stoichiometries of the defect search area were constrained to a subsection of the ternary plot.

All geometry optimisations in this work were performed with the PBESol XC functional (see table 9.6 for a full list of DFT parameters), which, while relatively fast, has a number of known shortcomings for modelling hBN. In 2018, Reimers *et al* compared DFT with a GGA XC functional, and with hybrid-GGA XC functionals (which includes a partial Hartree-Fock contribution) against more accurate post-DFT methods, such as coupled-cluster methods. Their results show that GGA and hybrid-GGA functionals are insufficient for accurately ranking excited states or for modelling optical spectra in defected hBN. No optical spectra, as modelled with the DFT parameters used in this work, have been reported in this thesis for this reason. The accuracy of any optical absorption spectra that can be modelled through CASTEP is improved upon by the modelling performed by Chen and Montserrat

with the GW approximation and BSE. Reimers *et al* report that when performing a geometry optimisation on a model compound (a finite planar molecule of hBN terminated with H) vacancies were often "healed", suggesting deficiencies in the accuracy of force predictions for GGA and hybrid-GGA XC functionals. However, they do not observe the same healing for geometry optimisations performed for a periodic system, such as the ones used in this work. This can be partly attributed to strain imposed on the system by the pristine hBN surrounding the defect and the fixed cell parameters.[148] Due to the high-throughput nature of this work, computationally costly hybrid-GGA XC functionals and post-DFT methods were not appropriate at the geometry optimisation stage. Although such methods would yield more accurate results than a GGA functional, the work of Reimers *et al* suggests that a GGA functional is sufficient for geometry optimisations in a 2D system and by compromising on accuracy for speed, a larger number of defects can be considered. Furthermore, this research aims to gain understanding the defects that appear in hBN not just through comparing the energies of defects, but by comparing the modelled optical spectra of these defects against the experimental spectra. Given this and that a optical spectra will be modelled for a broad range of defects (selected not just based on energy but also stoichiometry), there is a greater than typical tolerance for errors in energy rankings.

4.3 Defect Formation Energies

The formation energy, which is the energy change that occurs when a compound is formed from its constituent elements (see section 3.2 for more details), is an effective metric for comparing the stability of compounds with different stoichiometries. While formation energies may therefore seem practical for comparing defects, if formation energy was computed for an entire cell, it would be dominated by the contribution from the pristine hBN outside the defect. The implication of this is that the formation energy of the cell trends towards the formation energy of hBN as the number of pristine formula units of hBN around




Fig. 4.2 Plot summarising structures found through AIRSS. X-axis shows defect formation energy and Y-axis shows number of structures with a formation energy below that value. The figure shows that AIRSS favourably discovered relatively low-energy defects, with multiply-discovered structures appearing on the plot as column of points with a single formation energy. The discovered structures cover a broad energy range between 0 eV (pristine monolayer hBN) and 40 eV.

the defect increases. Instead, a defect energy, E_D , is defined as

$$E_D = E_T - E_{hBN}, \quad (4.1)$$

where E_T is the total energy of the cell and E_{hBN} is the energy of a pristine hBN cell of the same size. For a defect with the same number of atoms and the same stoichiometry as the pristine cell, E_D is equal to the defect formation energy, E_F . Otherwise, the formation energy may be expressed as

$$E_F = E_D - \sum_i n_i P_i, \quad (4.2)$$

where i is an arbitrary element, P_i is the chemical potential of that element, and n_i is the change in the number of atoms of that element with respect to the pristine cell. Which is to say that for an additional B atom, n_B is equal to +1, while for a B vacancy, n_B is equal to -1. Here, the chemical potential of B, N, and C, are taken to be the internal energy per atom of the elements in their ground state (a comparison of different chemical potentials is included in




Fig. 4.3 Ternary plot showing the distribution of $5 \times 5 \times 1$ hBN supercells generated during AIRSS. The position of each point represents the stoichiometry of the defect search area, not the whole cell. Each stoichiometry is coloured according to the number of structures generated with that stoichiometry. The vertices, B, C, and N, represent the a defect search area containing 100% B, C, or N, respectively.

the supplemental data as figure 9.2). When comparing the stabilities of different phases of a system, the formation energy is usually normalised by dividing by the number of atoms. This allows for the comparison of formation energy of two different phases with different numbers of atoms per cell. However, when calculating the defect formation energy per atom, E_{Fpa} , it is not appropriate to divide by the total number of atoms in the cell as this leads to E_{Fpa} becoming dependent on the size of the cell which contains the defect. Dividing by the number of defect sites (where a site may be vacancy, a substitution, or an interstitial) constituting the defect is also sometimes inappropriate as the number of sites constituting a defect may change. The $B_N C_B V_B$ defect, for example may rearrange to form $C_B V_N$ through migration of a B atom from a N site to a vacant B site. This rearrangement leads to a reduction in E_F with no change in the number of atoms, but the number of defect sites in the system has reduced

from 3 to 2, resulting in an apparent increase in E_{Fpa} . Normalising to the number of sites must therefore only be performed where the number of defected sites is unambiguous, such as when there are an irreducible number of defect sites. For example, C_B , and $C_B C_B$, in the absence of other neighbouring defects, may not rearrange to reduce the number of defected sites and therefore the difference in E_{Fpa} between C_B , and $C_B C_B$ provides meaningful insight about the relative stability of the two defects. Hence, the following formula is defined for an irreducible number of defect sites, nDS , per defect,

$$nDS = |\Delta C| + |\Delta B| + |\Delta N|, \quad (4.3)$$

where ΔC , ΔB , ΔN are the change in the number of C, B, and N atoms, with respect to the number of atoms in a pristine hBN cell of the same size. In the edge case where nDS is equal to zero, the atoms may rearrange to form pristine monolayer hBN: the ground state for this system. Under this condition, E_{Fpa} is taken to be equal to E_F , which is equal to E_D .

Given equations 4.1, 4.2, and 4.3, a normalised defect formation energy, E_{Fpa} , is therefore defined, which is independent of the size of the cell containing the defect and independent of the specific arrangement of the defect as,

$$E_{Fpa} = \frac{E_T - E_{hBN} - \sum_i n_i P_i}{|\Delta C| + |\Delta B| + |\Delta N|} \quad (4.4)$$

This formulation for a defect formation energy is sufficient for comparing the stability of two defects with the same stoichiometries, and is also useful for comparing the relative stability of two related defects, such as B_N and $B_N B_N$ but should not be used as a quantitative measure of the difference in stability between two defects with different stoichiometries, due to the ambiguity of defining the number of atoms with constitute a defect. For this reason, no defect convex hull is reported in this work.

4.4 Low Energy Structures

An uncertainty is estimated in the defect formation energies of 0.092 meV / atom from the k-point spacing and planewave cutoff selected for this work (see 9.6 for all DFT parameters). This uncertainty is dominated entirely by the planewave cutoff with an uncertainty of 0.091 meV / atom as compared to the 0.011 meV / atom from the k-point spacing, as shown in Figure 4.4. The uncertainty from the SCF random seed is estimated at 0.001 meV / atom and is therefore negligible. This uncertainty of 0.092 meV / atom is two to three orders of magnitude smaller than the difference in energy between defects; therefore the chosen k-point spacing and planewave cutoff are sufficient for comparing the stabilities of defects in monolayer hBN.




Fig. 4.4 Convergence of energy difference per atom between two hBN defects with the same stoichiometry with respect to k-point spacing and planewave cutoff. Y-axes are shared.

For all intrinsic defects and the lowest energy polymorph of each unique set of defect sites, an additional geometry optimisation has been performed for the defect in a $15 \times 15 \times 1$ cell. This allows for more accurate comparison of defect formation energies for the low energy defects by allowing for greater relief of strain over a 1412 \AA^2 area rather than

a 157 \AA^2 area. $15 \times 15 \times 1$ cells also allow for more accurate modelling of excitons and emission/absorption spectra by allowing for the excitons to be distributed over a larger area. The strain contribution to the total energy of the cell is reduced in the $15 \times 15 \times 1$ cells, and so the defect formation energy of a defect is expected to be smaller in a $15 \times 15 \times 1$ cell than in a $5 \times 5 \times 1$ cell. The mean change in formation energy was 0.136 eV , with most structures falling below this value. More significant changes were observed in cells with vacancies, which were more strained than similar defects without defects, and in the most extreme instance, the energy changed by 2.750 eV / atom . The distribution of this change is shown in Figure 4.5. Any increase in energy is attributed to insufficient k-point sampling. The k-point spacing selected for AIRSS was chosen as a compromise between accuracy and computational cost; while a finer k-point spacing would mitigate the error in energy from the k-point sampling, the memory demands were too large to practically perform geometry optimisations with more than one k-point for the $15 \times 15 \times 1$ cells. Any subsequent discussion of defect formation energies in this chapter refer to the defect formation energies for defects in a $15 \times 15 \times 1$ cell.




Fig. 4.5 Comparison of defect formation energy for defects in $5 \times 5 \times 1$ hBN supercell and $15 \times 15 \times 1$ hBN supercell. Most defects show a relatively small change in defect formation energy, $\sim 0.136 \text{ eV}$, between the $5 \times 5 \times 1$ and $15 \times 15 \times 1$ cells, while a small number of cells experience a difference in energy as large as 2.389 eV .

hBN defect modelling in the literature has largely focused on simple intrinsic defects. These defects are easy to manually construct, and the simplest of these defects, which are henceforth referred to as the “fundamental defects”, are useful for describing the other defects. With the constraint of three elements, B, N, and C, there are six fundamental defects, which are C_B , C_N , N_B , B_N , V_N , and V_B , in increasing order of defect formation energy. A sum of these energies and a visualisation of the defects are found in Figure 4.6. These values are largely consistent with the formation energies of fundamental defects in multilayer hBN reported by Weston *et al* in 2018. [149]. Weston *et al* report two sets of formation energies for a N-rich and N-poor environment, but their results show C_B to be the defect with the lowest formation energy, followed by C_N . The results for the remaining defects are more ambiguous with the ordering depending more significantly on the environment. N_B is shown as the next most stable defect in a N-rich environment with B_N very high in energy. In a N-poor environment, V_N is the lowest energy defect without C, followed by B_N , with N_B being higher in energy. However, N_B 's maximum formation energy is lower than B_N 's maximum formation energy, showing some consistency with the ordering reported in this work. In both Weston *et al*'s study and this study, V_N is shown to be more stable than V_B . There is a large difference in the relative stability of N_B/B_N and V_N/V_B between the two works, but this may be attributed to the neighbouring layers stabilising the vacancies: an effect which is absent in hBN monolayers. The lowest energy intrinsic defect was $C_B C_N$, in which the two C atoms are adjacent. The defect had a E_{Fpa} of 1.217 eV, showing a significant reduction in energy from the constituent defects, C_B and C_N which had a E_{Fpa} of 2.045 eV and 2.389 eV, respectively. The addition of another C atom to form a 3-C chain did not reduce the energy further but caused a small increase of only 58 meV in $C_B C_N C_B$ and 169 meV in $C_N C_B C_N$. Polymorphs of $C_B C_N C_B$ and $C_N C_B C_N$ were also discovered where the third C was disconnected from the chain and in this case the energies increased by an additional 221 meV and 238 meV, respectively. This indicates a stabilising effect from the formation of

graphite-like C-C bonds. A small energy reduction was also observed for pairs of C_B and C_N defects, which are $C_B C_B$ and $C_N C_N$. In these cases there was a reduction in E_{Fpa} of 38 meV and 34 meV, respectively, significantly less than the 1000 meV reduction for the formation of $C_B C_N$ from C_B and C_N , providing further evidence that the formation of C-C bonds is significant in stabilising the defects.

The substitution of C into B and N positions appears to be more favourable than the substitution of B into N positions and the substitution of N into B positions. The N_B and B_N defects had a E_{Fpa} of 3.210 eV and 3.808 eV, respectively. Vacancies in the B and N positions (V_N and V_B) were higher in energy still with a E_{Fpa} of 7.958 and 10.356. These general trends have been attributed to the ability for the hBN ring to distribute charge upon substitution. In the case of N_B , the additional N atom has an excess of two electrons which may be delocalised by donating charge to the neighbouring N atoms. Similarly, in B_N , the additional B atom has a deficit of two electrons which may be compensated for by withdrawing charge from the neighbouring B atoms. Substituting for a C atom instead of a B or N atom is more energetically favourable as this leads to a excess or benefit of only one electron, instead of two. The relatively low E_{Fpa} of $C_B C_N$ compared to its constituent defects may be attributed to the stabilising effect of the electron-deficient C atom withdrawing charge from the other C atom which is electron-rich. In the instances of $C_B C_N C_B$ and $C_N C_B C_N$, the charge is imbalanced and so the stabilising effect of the charge redistribution is reduced. Furthermore, it is observed that B-deficient defects are lower in E_{Fpa} than N-deficient defects, suggesting that N is better at accepting a partial negative charge than B is at accepting a partial positive charge.

So far, only intrinsic defects have been discussed, but AIRSS searches were also performed for extrinsic defects. The lowest energy defect identified, with a E_{Fpa} of 0.025 eV, was a N_2 molecule adsorbed to the surface of a pristine hBN monolayer. The distance between N_2 and the hBN was 3.25 Å, which is similar to the hBN interlayer distance of




Fig. 4.6 The defect formation energy per atom of the six fundamental defects. The defect search areas have been inset with N shown in blue, B in green, and C in brown. Normalised formation energies calculated from chemical potentials of B, N, and C in their ground states.

3.50 Å. The second lowest energy defect was similar and took the form of a cyanogen-like C_2N_2 molecule, with alternating C-N bonds, also adsorbed to the surface of pristine hBN. This defect had a E_{Fpa} of 1.008 eV. In addition, several low E_{Fpa} defects were identified that took the form of linear chains of B, N, and C atoms perpendicular to the plane of the monolayer. These defects bond to the monolayer at a single site, minimising the impact on the structure of the hBN. However, at the point of contact, the atom in the plane of the monolayer undergoes sp^3 hybridisation, rather than sp^2 hybridisation seen elsewhere in hBN. This is associated with a shift from a trigonal planar geometry to a tetrahedral one, which imposes strain on the monolayer. Despite this, several of these defects were identified that were lower in E_{Fpa} than the fundamental defects, suggesting that the addition of atoms to the surface of the monolayer is energetically more favourable than disrupting its structure. A summary of some of the low energy axial defects can be found in Figure 4.7.




Fig. 4.7 Visualisation of a sample of low energy axial defects and their normalised defect formation energies. Formation energies calculated from chemical potentials of B and C in their ground states. The chemical potential of N is taken as the difference between the energy per formula unit of hBN monolayers and the chemical potential of B. Z-axis given by blue arrow with the XY plane perpendicular to arrow. Pristine hBN surrounding defect not shown to improve visual clarity. The whole cells as viewed from the Z-axis can be found in figure 9.3.

4.5 Conclusion

I have applied AIRSS to generate over 5000 cells containing defects in monolayer hBN and compared the defect formation energy per atom of each defect. Initial defect searching was performed in $5 \times 5 \times 1$ hBN supercells and defects with defect formation energies less than 1 eV / atom were also geometry optimised in $15 \times 15 \times 1$ hBN supercells improve the accuracy of modelled optical absorption/emission spectra, and to minimise strain for more accurate energy comparisons. It was also discovered that C-rich defects were lower in energy than both N-rich and B-rich defects, and B-deficient defects were lower in energy than N-deficient defects. Further, substituting either a B or N atom adjacent to a defect with a C atom reduced the E_{Fpa} , suggesting a stabilising effect from the C. These general relationships have been attributed to the monolayer's ability to redistribute a change in electron density. Located between B and N on the periodic table, substituting B or N for C has a smaller

net change in charge density than substituting B for N or N for B. In addition, it is shown that axial defects are relatively stable in monolayer hBN. These defects have generally been overlooked in the literature thus far but may play an important role in single-photon emission in monolayer hBN. Further geometry optimisations will be performed on 15x15x1 cells to generate as many defects as possible, and the lowest energy defects will be provided to Mr. Chen for modelling optical absorption / emission spectra. The resultant spectra will be compared to experimentally observed single-photon emissions so that the emission spectra may be attributed to specific defects. Further, this may allow for the identification of defects that do not occur naturally during hBN synthesis but may act as SPEs of photons with useful frequencies for various purposes. Knowledge of the specific defects that give rise to single-photon emission will aid the design and engineering of hBN SPEs in the future.

Chapter 5

Potassium Graphite Intercalation Compounds

5.1 Introduction

Graphite-intercalation compounds (GICs) were first synthesised in the early 20th century [150, 151] but had little practical application before the introduction of Sony's LiCoO_2 - LiC_6 Li-ion battery [152, 153]. Li has the smallest ionic radius and the lowest atomic mass of the alkali-metals, making it highly mobile and allowing it to transfer more charge per unit mass than the other alkali metals. It is consequently an obvious choice as a charge carrier for a battery, but its abundance in the earth's crust is low (0.002%), compared to Na (2.36%) and K (2.09%) [154], making it relatively challenging and expensive to extract. Na-ion batteries have been studied extensively with a range of materials including transition metal, group 14, group 15, and group 16 electrodes, as well as graphite-intercalation electrodes [155–161]. However, poor interaction between Na^+ and graphene layers [158] leads to a low ratio of intercalation and an inherently poor charge capacity compared to the Li^+ analogue. Although K is at a disadvantage to Na in both abundance on earth and ionic mass, K^+ can intercalate to graphite at a ratio of 8:1 [162] (only 15% less efficient than Li^+). KC_8 anodes may offer

more charge capacity per unit currency than LiC_6 as GICs are generally cheap and easy to synthesise. Additionally, the benefits of furthering the understanding of KC_8 are not limited to battery applications as KC_8 has also been demonstrated to exhibit low temperature superconductivity [163–167, 167]. Synthesis occurs in a step-wise process through a series of intermediates termed “stages”. Stages 1-5 K-intercalated GIC (KGIC) within the Rüdorff-Hofmann model [168] are shown in Figure 5.1. The saturated GIC is designated “stage 1” and the numeral of the stage relates directly to the number of graphene layers between each intercalant layer. These materials have been synthesised via a gas-vapour [162, 169, 170] method historically, but electrochemical methods, which offer better rate control and allow for the reaction to be paused and reversed, have now become more common [171–173].

Though in some cases it is possible to observe the progress of intercalation with the naked eye (KC_8 has a distinctive golden colour and KC_{24} is blue), Raman spectroscopy offers a reliable, precise, and unobtrusive method of assessing the progression of the intercalation [174, 170, 163, 175]. KGICs are characterised by the Fano peak [176] at $\sim 1400\text{ cm}^{-1}$ [170], the G peak at $\sim 1550\text{ cm}^{-1}$, associated with in-plane C vibrations, and the 2D peak at $\sim 2700\text{ cm}^{-1}$, which is the overtone of the D peak at $\sim 1350\text{ cm}^{-1}$ [175]. Although the 2D peak is prominent in both graphite and graphene, the intensity of the peak diminishes during intercalation; allowing for it to be used as an approximate measure of intercalation [175, 163, 177, 178]. The G peak is of particular interest as it undergoes the most significant changes between each stage of intercalation [179, 175, 163]. Pristine graphite displays a Raman G peak at 1580 cm^{-1} [180] which splits into a doublet upon addition of an intercalant [170]. These two peaks are termed the “interior” and “boundary” peaks (henceforth referred to as G_i and G_b) and are associated with in-plane vibrations of the interior and boundary graphene layers, respectively. In the literature, a graphene layer is considered to be interior if it is bounded on both sides by other graphene layers, and a graphene layer is considered to be boundary if it is bounded on one side by K^+ and on the other side by a graphene

layer [181, 182, 163]. As the relative intensity of G peaks is related to the ratio of interior layers to boundary layers, it has been proposed that this could be used as a direct means of determining the stage of intercalation [163], however, this is not straightforward as the relative intensity of the peaks is influenced by the frequency of the laser [174]. The dependence of the intensity on the ratio of the layers is most apparent with stage 2 KGIC (KC_{24}) and stage 1 KGIC [174, 170, 163, 175]. Stage 2 features only boundary layers and therefore the Raman spectrum expresses only one G peak. Stage 1 features only graphene layers which are bounded by K^+ on both sides and the Raman spectrum expresses only a shoulder, or shallow peak, in place of the sharp G peak. Because of the relationship between the ratio of interior and boundary layers and the ratio of the intensity of the G peaks, interior and boundary have been inherited as labels for each of the peaks. Although interior and boundary are appropriate descriptors for the different layers in GICs, they may be misleading as labels for the two G peaks.

Density-functional theory (DFT) has been used for modelling GICs in the past. This includes structure characterisation [164, 183, 161], bandstructure and electronic density of states calculations [184, 170, 163], and phonon calculations [170, 185, 163, 186]. Though the latter has largely focused on the Raman spectra. In this work, DFT was used to model the complete phonon spectra for stage 1-5 KGICs and pristine graphite. The evolution of the phonon G peak is contrasted against the evolution of the density of states to highlight the relationship between bond strength and vibrational frequency. Two major models exist for the structure of KGICs. The earliest and most widely studied model is the Rüdorff-Hoffmann (RH) model [168], which features alternating layers of graphite and K^+ layers where the number of sequential graphite layers scales linearly with the numeral of the stage. In the RH model, each of the layers are planar and for an arbitrary XY plane through the structure, either only C atoms or only K^+ ions can be found. However, in 1969, Daumas and Hérold [187] proposed a new model in which the intercalants do not appear homogeneously distributed in

discrete planes but instead form “islands” of intercalant clusters (henceforth called domains) in arbitrary planes (a schematic of RH and DH unit cells can be found in figure 5.2). KGIC cells with both models were modelled and compared. A summary of DFT parameters and settings used for this project is found in Table 9.4. All RH structures have been manually constructed based on schematics available in the literature [168, 162, 164, 165, 15, 156], and then geometry optimised. Stages 1-5 RH and DH cells have been constructed such that they are charge neutral and satisfy the stoichiometries KC_8 , KC_{24} , KC_{36} , KC_{48} , and KC_{60} . Details about DH cells can be found in section 5.2. This research is supported by a collaboration with Prof. Mauro Pasta and Dr. Isaac Capone at the University of Oxford who provided experimental Raman spectra for the electrochemical intercalation of K into graphite.




Fig. 5.1 Visualisation of optimised Rüdorff-Hoffmann (RH) unit cells for stages 1 - 5, which have been used for modelling phonon spectra and electron density of states for cells within the RH model. C atoms shown in brown and K in purple. The numeral of the stage can be seen to match the number of C layers between each K layer.




Fig. 5.2 Schematic of unit cells for stage 2 KGIC with the RH (left) and DH (right) model. C layers are represented in brown and K ions in purple. The left schematic shows flat graphene-like layers where the spacing between layers is greater when separated by K ions than when not. The right schematic shows that all layers are bounded by K ions, but each interlayer spacing is less saturated. C layers are curved around K domains such that the interlayer distance is large across K ions and small across empty space. The XY plane is in plane with the RH layers and the Z axis is perpendicular to that plane.

5.2 Daumas-Hérolde Defects

5.2.1 Introduction

The distribution of the intercalants in GICs has a profound impact on the graphite layers. In the RH model for KGICs, interior layers have an interlayer distance comparable to graphite, but where two graphite layers are separated by K ions, the distance is enlarged in order to accommodate them. In the DH model, there is only one type of graphite layer and so all layers must be able to accommodate the intercalants. One may therefore expect graphite layers to be uniformly distributed at the minimum distance that can accommodate the intercalant. Instead it is observed that the graphite layers bend around the intercalants such that the interlayer distance across intercalants approaches that of the RH model, while interlayer distances elsewhere approach that of graphite. Several factors play a role in determining the gradient of these bends. The obvious are the preferences for large inter-layer distances across an intercalant and small inter-layer distances across empty space. However, there is also a preference to maintain optimal C-C bond lengths and C-C-C bond angles. π -conjugation

plays a significant role in stabilising graphene and graphite, and bends in the graphite layers disrupt this conjugation. Hence, it is energetically favourable to minimise the gradient of the graphite layers. Additionally, a bend in a layer may occur in proximity to a bend in a neighbouring layer, and because neighbouring layers interact, the gradient of one layer may be affected by the gradient of another.

The stacking of the layers is also impacted. Graphite favours AB stacking, as do interior layers in the RH model, but neighbouring boundary layers favour AA stacking so that the intercalant can be accommodated between two C-hexagons. The RH model is well-defined and when constructing model cells for simulation, the only true unknowns are the intraplanar C-C bond distance and the interplanar C-C bond distances (which for boundary layers is affected by C-K bond distances). These variables are all impacted by the PSP, XC functionals, and SEDCs applied, and can be optimised through a geometry optimisation. DH cells are not well defined as the positions of the K ions are unknown. In addition to all the variables listed for RH, one must also be concerned with the number of K ions per domain, and how those domains are distributed through the overall structure. Both of which also have an impact on the bends in the graphite layers, which are discussed in further detail in the preceding paragraph. Further to this, each graphite layer in a DH cell is only periodically in proximity to a K ion, which raises the question as to whether the graphite layers favour AB stacking or AA stacking.

There has so far been no experimental measurement of domain size in KGICs or any other GICs, nor are there any computational studies which make strong claims about domain size for different stages. However, in their book, Dresselhaus and Dresselhaus [188] claim that the staging behaviour in RH compounds arises in order to minimise lattice strain in the Z-axis; that is to say, it is energetically unfavourable to increase the distance between graphite layers, and so K ions localise to as few layers as possible to maximise the number of graphite layers with optimal interlayer spacing. They go on to argue that this strain is also a driving force

for the formation of large domains in DH cells, based on a study of molecular alignments in molecular crystals by de Gennes and Prost [189]. It is clear, however, that the clustering of ions in a DH cell would have no impact on the length of the Z-axis as all layers are occupied and therefore all neighbouring graphite layers experience an increased interlayer spacing at some point along their length, regardless of the size of the domain. However, the bending of layers around domains reduces the overlap coefficient between p_z orbitals on neighbouring C atoms, weakening π -conjugation and imposing strain on the layers. Increasing the size of the domain minimises the number of bends needed to accommodate the intercalant, minimising the strain on the material, as domain size trends to infinity, the material approaches an approximation of the RH model. The size of the domains is likely to be inversely proportional to the stage index in order to accommodate the slipping of the graphite layers to AA stacking that is required to accommodate the intercalant [188]. Furthermore, during intercalation, it is energetically favourable to accommodate new intercalants in an existing domain rather than creating a new domain as this minimises curvature in the graphite layers and hence minimises strain.

Given the number of unknown variables for DH cells, modelling is therefore firstly a structure-searching exercise.

5.2.2 Structure Searching Method

AIRSS is not an appropriate method for searching for DH cells. Defining DH cells requires hundreds of atoms and fully random structure searching will rarely yield DH cells. The large number of atoms and P1 symmetry of DH cells naturally leads to a large phase space. Instead, because the general features of the material are known, a more efficient direct searching method was applied. The first geometry optimisation was performed for a stage 4 cell, because the optimised stage 4 RH cell could be manually converted into a candidate stage 4 DH cell more directly than any of the other RH cells could be converted into their DH

analogues. The stage 4 DH candidate cell used the same cell parameters as the optimised stage 4 RH cell, except for the length of the cell axis perpendicular (The Z-axis) to the C-planes. This ensured that bond lengths and angles in the graphite layers were unchanged. As the interlayer distances and the bending in layers was unknown, the distance between layers was set to the midpoint between the boundary interlayer distance and the interior interlayer distance in the RH cells, and no bending was present in the layers so as to avoid blindly biasing the geometry optimisation. The candidate cell was then geometry optimised to a force and stress tolerance of $0.01 \text{ eV } \text{\AA}^{-1}$ and $0.01 \text{ eV } \text{\AA}^{-3}$. AA stacking was deemed to be the most likely stacking arrangement in order to better facilitate K ions and therefore this stacking was chosen for the candidate cell. Figure 5.3 shows the manually-constructed candidate stage 4 DH cell and also the cell after it has been geometry optimised.




Fig. 5.3 Candidate stage 4 DH cell before geometry optimisation (left) and after geometry optimisation (right). The bends in the layers are discovered naturally by the geometry optimisation algorithm. The length of the cell in the X and Y directions was largely unchanged, however, the Z length increased from 13.3 \AA to 16.6 \AA .

Although the process of generating the first DH cell was successful, it was extremely computationally expensive. By default, CASTEP only computes k-points for an irreducible representation of the BZ, to minimise computational redundancy and thereby minimise compute time. However, as DH cells have P1 symmetry, it is necessary to compute all k-points, making the calculation more expensive than for other KGIC cells of similar volume / composition. The geometry optimisation also seemed to require more iterations before

the forces and stresses reached acceptable tolerances; suggesting that the potential energy landscape may be more complex for DH cells than RH cells.

The number of atoms required to define a cell increases with the size of the domain and DH cells with domains as small as two ions are prohibitively expensive to model with DFT. Consequently, all DH cells are explored with a domain of a single ion, which minimises cost and also sets an upper bound for the energy of DH cells. DH unit cells are sometimes defined with a number of graphite layers equal to the stage index. These structures minimise interactions between K in the Z-axis but even with only two graphite layers, periodic images are no less than 8 Å apart. This research focuses on 2-layer cells and domain sizes of one because it was computationally too expensive to model phonon spectra for cells with more layers and larger domains.

In order to minimise the number of geometry optimisation iterations, the starting cell should be as similar as possible to the optimised cell. To this end, the fully-geometry optimised stage 4 DH cell is used as a template for the construction of additional cells. Firstly, constraints were defined. As mentioned in the preceding paragraph, domains are constrained to only one K ion, and therefore only one K ion is included per layer. This necessitates that each graphite layer should contain a number of C atoms equal to the ratio of C:K of the intercalation stage. Additionally, a minimum of two layers is required per cell in order to facilitate bends in the graphite layers. As the cell parameters for the cell are unknown, several different cell parameters were tested, which were constructed from two different “base” cells containing two and four atoms, see Figure 5.4. For each stage, multiple graphite supercells are then constructed from these base cells such that they contain the desired number of layers, and C atoms in each layer, and such that each supercell has a different X:Y length ratio. Each layer was spaced apart by 4.15 Å. For each supercell, one K ion is then inserted in a fixed position in the first layer and all combinations of K ions in each interstitial site in the remaining layer(s) are generated. From the calculations performed on RH cells, it is known

that charge is donated from K ions to neighbouring C atoms. In order to best distribute this charge across the graphite layers, the K ions should themselves be evenly distributed across the cell. For each combination of occupied interstitial sites, the K-K distances are computed and the cell with the greatest K-K distances is selected in order to best distribute the charge. The distance between each C atom and each K atom was then computed, and the value of the C atom's Z-coordinate was shifted according to some polynomial that has been fit to the bends discovered in the original stage 4 DH cell.




Fig. 5.4 2-atom hexagonal (left) and 4-atom orthorhombic (right) base cells used to construct other DH cells. C atoms are shown in brown. The X and Y axes are equal in length in the hexagonal base cell but not in the orthorhombic cell. Consequently, for arbitrary cell lengths A and B, hexagonal cells of dimensions (A,B) and (B,A) are degenerate but orthorhombic cell of dimensions (A,B) and (B,A) are not.

Each cell is then geometry optimised to sufficient force (1×10^{-2} eV Å⁻¹) and stress (5×10^{-2} GPa) tolerances. In order to minimise the cost of geometry optimisation steps, an initial geometry optimisation was performed with an MLP but with fixed cell parameters. Training the MLP was possible because of the DFT geometry optimisations that have already been performed on RH cells and the candidate stage 4 DH cell. Every step from these geometry optimisation was extracted to construct a database of approximately 2100 cells. However, this database is compromised of effectively only eight (unique) structures (RH stages 1-5, DH stage 4, and graphite). The contents of the database are consequently highly correlated. Because of this, the ANN was trained on a relatively small fraction of the database

(400 structures) in order to minimise overfitting and the computational cost of training a model. The remaining 1700 structures were implemented as a validation set. In some cases, it is desirable to localise all steps from a geometry optimisation walk to either the training, validation, or test (if included) set so as not to bias the validation error. Particularly towards the end of a geometry optimisation, the difference between one step and the next may be smaller than the sensitivity of the ANN to atomic positions, and therefore splitting two neighbouring steps between the training set and the validation set is almost analogous to including a single structure in both the training set and the validation set. Given the small number of geometry optimisations included in the database, structures were assigned to the training and validation set randomly, but this comes with the caveat that, due to overfitting, the errors in the validation set may be somewhat smaller than the errors on force predictions for any new DH cells.

This training and validation set was used to train an MLP with the PaiNN [76] architecture, as implemented in SchNetPack [77, 78]. Specific details can be found in Table 9.5. Over 4000 training steps, an ANN was trained with a force MAE of $3.6 \text{ meV } \text{\AA}^{-1}$ and $6.0 \text{ meV } \text{\AA}^{-1}$ for the training set and validation set, respectively. The progression of force errors over the course of training is shown in Figure 5.5 and a comparison of the ANN-predicted forces and DFT forces for the validation set are given in Figure 5.6. The ANN was trained exclusively on forces because the energies of the structures at this stage of structure searching would not be compared (given that a geometry optimisation would later be performed with DFT). The total energy of the cell is used to estimate the Hessian in the geometry optimisation process (see section 2.1.2) but if the ANN correctly describes the shape of the PES, the exact value of the energy is unimportant so long as the energy prediction is internally consistent. It was desirable to also train an ANN that could predict stresses, either as a separate ANN or as part of a ANN that can also predict forces. However, even though an ANN could be trained with seemingly sufficiently small stress errors, the ANN proved unsuccessful on practical

implementation as the geometry optimisations often lead to cells becoming unrealistically large or small. However, using an MLP that was trained on only forces, it was possible to perform accurate geometry optimisations with fixed cell parameters.




Fig. 5.5 Plot to show the convergence of the force errors for the training and validation set over the course of training.

The performance of the MLP was tested against stage 2 DH cells. Unoptimised structures generated via the method described in the preceding paragraphs each underwent a fixed-cell geometry optimisation independently with DFT and with the MLP. A DFT singlepoint calculation was then performed on all structures. In all cases, the difference in energies between the DFT-optimised and the MLP-optimised structures was less than 0.15 meV per atom. The success of this approach is likely due to the limited phase space that the ANN is required to learn, and also the similarity between the unoptimised algorithm-generated structures and the optimised structures. Furthermore, the MLP geometry optimisation reliably reproduced the DFT optimised structure even if the starting structure started with flat C-layers only. However, this was more computationally demanding than estimating the bends in the C-layers with a polynomial and geometry optimising from there. In order to optimise the cell parameters, after optimising with the MLP, a geometry optimisation was performed with DFT using the same parameters used to optimise the RH cells.




Fig. 5.6 Plot to show ANN predicted forces against DFT forces for validation set. Axes have been constrained to $\pm 10 \text{ eV } \text{\AA}^{-1}$ to improve the readability of the plot as a negligible number of points exist outside this range. The MAE reported on the figure reflects this constrained force range, rather than the total validation set.

This method can be applied to cells of any intercalation stage, containing domains of any size, and containing any number of layers. Large cells such as in Figure 5.8 are too computationally demanding to geometry optimise with DFT but the structure was optimised with an MLP within minutes on a single core.

5.2.3 Structure Analysis

The stability of the various cell representations are assessed by comparing the total energy per atom. The stoichiometry for a given stage is fixed and therefore the difference in formation energy between different cells of the same stage is the same as the difference in total energy. The energy of all cell representations considered within this research are relatively small, with the most significant difference being $\sim 59 \text{ meV/atom}$ and occurring within stage 2 between the RH cell and a $1 \times 12 \times 2$ hexagonal DH cell. However, this DH cell was particularly high in energy compare to the other cells, likely due to the X-axis length of 2.5 \AA . This resulted in K ions occupying neighbouring hexagons, which is never observed with RH cells. Cells




Fig. 5.7 Example optimised DH cells for stages 2 - 5. C atoms shown in brown and K in purple. Raman spectra for these cells can be found in 5.14. Note that the graphite layers in stage 2 are almost planar due to the packing density of the K ions.

were generally lower in energy when the axes were of a similar length, which is attributed to improved flexibility in the curvature of the layers and more evenly spaced K ions. It is difficult, however, to draw specific insight into the relationship between cell geometry and energy. For stage 5 orthorhombic cells, the $3 \times 5 \times 2$ cell was 20 meV/atom higher in energy than the $5 \times 3 \times 2$ cell, suggesting that it is favourable for the zig-zag edge to be longer than the armchair edge, but for stage 4 orthorhombic cells, $4 \times 3 \times 2$ was 6 meV/atom higher in energy than $3 \times 4 \times 2$, which would imply the opposite. A more extreme example appears for hexagonal cells. At stage 5, the $3 \times 10 \times 2$ cell is 12 meV/atom higher in energy than the $2 \times 15 \times 2$ cell, but for stage 3, the $2 \times 9 \times 2$ cell was 8.5 meV/atom higher in energy than the $3 \times 6 \times 2$ cell. Narrow cells, such as stage 3 $2 \times 9 \times 2$ feature a minimum K-K similar to that observed for RH cells, ~ 5.0 Å. In these cells, no curvature is observed along the zigzag edge meaning that the cells effectively contain infinitely-long 1-dimensional K domains. It is expected that cells with larger domains have lower energy as curvature in the C-layers




Fig. 5.8 The structure searching approach used for DH cells is scalable to cells of any number of layers and with domains of any number of ions. This cell contains 784 atoms, and is consequently too computationally expensive to perform a phonon calculation, and it is impractical to perform a DFT geometry optimisation. However, it was possible to perform a geometry optimisation with the MLP in less than 17 minutes.

is minimised, but when cells included infinitely-long 1-dimensional K domains, the energy was relatively high. Infinitely-long 1-dimensional K domain may paradoxically increase the strain in the system, by inducing an infinitely-long 1-dimensional sharp kink in the C-layer along the K domain, as opposed to the more gradual curvature around single-ion domains. For very large 2-dimensional finite domains, the energy penalty associated with this kink is outweighed by the large flat area C-layers above and below each domain.

One energy-geometry relationship has expressed itself clearly. When the initial cells were generated with only the Python script described in preceding section, graphite layers were uniformly curved around the K ions, however, when the cells were subject to a MLP or DFT geometry optimisation, the extent of the curvature in the graphite layers changed depending on the length of the cell in each direction. Due to periodic boundary conditions and the single K ion per layer, the length of each of the axes determined the distance between K ions within a layer in each direction. Curvature is more significant in directions in which the ions are far apart. No curvature was observed for ions spaced less than 5 \AA apart and at $\sim 7.5 \text{ \AA}$,

the Z-coordinate of the C atoms differed by less than 0.1 Å. The longest orthorhombic cell generated had a cell length of 21 Å and the Z-coordinate of the C atoms differed in that direction by as much as 1.7 Å.

All DH cells generated in this work were higher in energy than the RH cells by 12-24 meV / atom. DH KGICs with domain sizes as small as one atom are therefore likely to be energetically accessible at ambient temperature. A summary of the 2-layer DH cells with their relative energies and layer distortion can be found in Table 5.1.

5.3 Electronic Structure

KGICs are structurally similar to graphite however, while graphite displays AB stacking throughout, KGICs display AA stacking where would-be neighbouring layers are separated by a K layer. To this extent, KGICs can be treated as a graphite framework with weakly interacting K ions between some (or all) of the layers. These K ions donate a small amount of charge to the neighbouring C layers, resulting in negatively charged C layers and positively charged K layers. Mulliken analysis of RH cells (see Table 5.2) shows that charge on the C atoms increases with K-intercalation and that the charge transfer is greatest on C atoms that neighbour K ions (the boundary layers). The electronic structure of KGICs are not unlike that of graphite or graphene (with stage 1 being closer to graphene than graphite) and consequently, they display similar electronic properties. Accommodating K ions increases the C inter-layer distance such that the C layers begin to behave more like non-interacting graphene layers than graphite layers, leading to a bandstructure that very closely resembles graphene. DFT is well known for underestimating band gaps in semiconductors; particularly with the LDA XC functional, but also with a GGA XC functional such as the one used here. The model appears to accurately reproduce the graphene Dirac cone, but it is likely that the bandgap in the Z-direction is underestimated for all structures considered in this work. Although graphene is a zero-band gap semiconductor, K-C charge transfer raises

the Fermi-level from the center of the Dirac cone to what otherwise would have been the conduction band. For this reason KGICs are 2-dimensional metals with a positive relationship between conductivity and K concentration [190, 162].




Fig. 5.9 Electronic density of states of stage 1-5 RH KGIC and graphite computed with DFT. C s states are shown in light brown, C p states in dark brown, K s states in light purple, and K p states in dark purple. K-intercalation has a minor impact on the shape of the C states but the Fermi level can be seen shifting to higher energies at greater levels of intercalation. Figure 9.1 shows an alternative representation where the DOS is zeroed to the Dirac cone. Details of the DFT parameters used can be found in 9.4

The most relevant features of the band structures are best summarised by electronic density of states (DOS) plots. Figure 5.9 shows the electronic DOS for stages 1-5 RH KGIC and graphite, with the energy zeroed to the Fermi energy of each compound. Although the C contributions to the electronic DOS are largely unchanged by K intercalation, there is a significant increase in K contribution above the Fermi-level. The ratio of the C and the K contribution to the overall electronic DOS scales with the ratio of C atoms and K ions in the compound. In addition, charge transfer from the K ions is responsible for the climbing Fermi-level and the occupation of molecular orbitals that otherwise would have been unoccupied. In the case of stages 2-5, non-bonding orbitals are occupied. However, at stage 1, anti-bonding orbitals are occupied, weakening the C-C bonds and reducing the

vibrational frequency. Furthermore, the Fermi level rising to an energy with a higher DOS is responsible for the increase in electronic conductivity.

Electronic DOSs have also been computed for the lowest energy DH structures, Figure 5.10, so that they can be compared with the DOSs of the RH cells. Evolution of the DH DOSs displayed similar behaviour to the evolution of the RH DOSs, with the general shape of the DOS being largely unchanged from graphene, but with the Fermi-level climbing above the graphene Dirac cone and the K s and K p states appear above the Fermi-level. The same limitations when modelling Raman spectra of RH cells therefore also apply to modelling Raman spectra of DH cells. Furthermore, the same inferences are made about the relationship between the DH phonon spectra and the electronic DOS as made in section 5.3.




Fig. 5.10 Electronic density of states of stage 2-5 DH KGIC computed with DFT. C s states are shown in light brown, C p states in dark brown, K s states in light purple, and K p states in dark purple. The shape of the DOS does not differ significantly from the analogous DOSs for the RH cells. See Figure 5.9. Corresponding DFT parameters are reported in table 9.4

5.4 Phonon and Raman Spectra

Finite displacement phonon calculations were performed for every RH KGIC cell generated as part of this research, as well as for graphite and the low-energy DH KGIC cells. Raman calculations for KGICs were not possible as the materials are low-dimensional metals. Instead

simple phonon calculations are performed and the space group of the phonon eigenvectors are used to determine whether a mode is Raman active. This approach gives a qualitative Raman spectrum but lacks any information about the intensity of the Raman modes. Lorentzian broadening has been applied to many of the computational phonon and Raman spectra reported in this study, in order to model the broadening observed experimentally. The intensity of these broadened peaks is derived from the number of phonon modes “inside” the peak, rather than the Raman intensity of those modes.

Most of the phonon modes identified in stages 1-5 RH cells were assessed to be Raman-active, as shown by the difference between Figures 5.11 and 5.12. The phonon spectra, by definition, contain all the modes from the Raman spectra as well as a number of Raman-inactive modes. The difference in activity is most pronounced for the graphite phonon modes in the range $1300\text{--}1400\text{ cm}^{-1}$, and for stage 1 (C_8K), which has three Raman-inactive peaks below 1100 cm^{-1} . To verify the accuracy of the models, the modelled spectra are compared to the experimental Raman spectra collected by Dr. Capone, reported in Figure 5.13. These measurements were made with a 532 nm laser during the electrochemical intercalation of K^+ into graphite. Spectrum 1 therefore shows the Raman spectrum of graphite and spectrum 11 is attributed to the spectrum for the saturated KGIC, C_8K . The intermediate spectra are more challenging to attribute to specific stages as the changes in the spectra are gradual. However, spectrum 6 contains a single peak, which is conventionally attributed to stage 2 (KC_{24}) as the stage 2 RH cell contains only one type of C layer.

Multiple differences can be observed between the modelled spectra, and the experimental spectra. As the Raman spectra collected by Dr. Capone are largely consistent with the experimental spectra reported elsewhere in the literature, the differences are largely attributed to shortcomings in the model. Differences in the spectra between 1300 and 1400 cm^{-1} were expected due to the use of the finite displacement phonon method. Experimental Raman measurements taken with a 785 nm wavelength laser (see the pink spectra in Figures 5.11




Fig. 5.11 Complete phonon spectra for RH KGICs and graphite. Experimental Raman spectra for Stage 1 KGIC with 512 nm and 785 nm lasers are superimposed over the DFT stage 1 phonon spectrum for comparative purposes. Spectra should only be compared qualitatively, not quantitatively as intensities are normalised to the maximum of each respective spectrum. Corresponding DFT parameters are reported in table 9.4.

and 5.12) show a broad Fano peak in this region for all spectra. Finite displacement does not consider resonant Raman effects and therefore cannot reproduce the Fano peaks. Instead, a large number of similar-energy phonon modes were observed in this range across all spectra. Interestingly, these modes mostly show as Raman active for stages 1-5, but all appear Raman-inactive for graphite. Isaac Capone's 785 nm spectra also show a peak at 750 nm and 900 nm, which are constant across the spectra. All stages show a peak which is associated with the same phonon mode, vibration of the C layers in the Z-direction. For stage 1 this occurs at 805 cm^{-1} and for graphite it occurs at 865 cm^{-1} , while the intermediate stages show the peak at around 855 cm^{-1} but in the presence of other phonon modes with a similar energy that are not present for stage 1 and graphite in that area of the spectrum. Stages 2 to 5 also show consistent peaks at 987 cm^{-1} , which is associated with in-plane twisting of the hexagons,




Fig. 5.12 Raman spectra for RH KGICs and graphite. Experimental Raman spectra for Stage 1 KGIC with 512 nm and 785 nm lasers are superimposed over the DFT stage 1 phonon spectrum for comparative purposes. Spectra should only be compared qualitatively, not quantitatively as intensities are normalised to the maximum of each respective spectrum. Corresponding DFT parameters are reported in table 9.4

and at 1201 cm^{-1} , which is associated with in-plane movement of half the C-atoms. The modelled spectra and experimental spectra all show a downshift of the phonon modes as the stage of intercalation increases, but this is more significant in the modelled spectra. This downshift is attributed to the charge transfer from the K ions to an anti-bonding HOMO, which destabilises the C-C bonds. The most significant discrepancy occurs at the G-peak, which is shown in greater detail in Figure 5.14. Experimentally, a single peak is seen below 1600 cm^{-1} for graphite and as K is intercalated, an additional peak grows in intensity 30 cm^{-1} higher in wavenumber: these two peaks make up G_i and G_b , respectively. As the stage index decreases, the intensity of the interior peak diminishes, while the higher wavenumber boundary peak grows until G_i disappears entirely at stage 2. Experimental spectra for stage 1 typically show no peak or a broad shoulder. While the computational Raman G-peak is




Fig. 5.13 Evolution of experimental Raman spectrum during electrochemical intercalation of K^+ into graphite. Indices on Y-axis show sequential measurements of Raman spectra made at regular time intervals. Graphite is shown at the bottom of the figure, and stage 1 is shown at the top of the figure. Raman measurements were made with a 532 nm wavelength laser. Spectra intensities are normalised to the maximum of each spectra and so intensities should not be compared across measurements. Figure provided by Dr. Capone and reproduced with permission.

consistent with experimental spectra for stage 1 and graphite, the Raman spectra for the intermediate stages with both the RH and DH models are inconsistent with experiment. The DH spectra diverge significantly from experiment with a large number of Raman modes appearing across the range $1500\text{--}1600\text{ cm}^{-1}$. The Raman spectrum of each DH cell generated in this project is complex and were sensitive to the shape of the cell. The spectra included in Figure 5.14 are those found during the structure searching process which were expected to best reflect reality, see section 5.2. However, the spectra do not differ significantly within a

stage. The DH model, and in particular, the DH model with single-ion domains therefore do not seem to be an accurate representation of KGICs synthesised electrochemically.

The apparent intensities of the peaks in Figures 5.11, 5.12, and on the right hand side of figure 5.14 arise entirely from the Lorentzian broadening that has been applied to reflect the broadening that appears experimentally. Because of this, the height of the peaks in these figures reflects the number of phonon modes within a given wavenumber range. Although this is a contributing factor to peak intensity in experimental spectra, the change in polarisability is also a significant factor and there is no known method for modelling this in metals. The Lorentzian intensity estimate is likely entirely inadequate and the Raman spectra with Lorentzian broadening are therefore misleading. The experimental spectra is subsequently compared to the computational spectra *without* broadening, as shown on the left side of Figure 5.14. The spectra for stages 2-5 appear to be comprised of two clusters of phonon modes, with a splitting of 30 cm^{-1} , which is entirely consistent with experiment. Raman modes at $\sim 1530\text{ cm}^{-1}$ for stages 4 and 5 may be interpreted as two peaks instead of one, however they are only split by $\sim 10\text{ cm}^{-1}$, which is similar to the FWHM maximum of the experimental peak, and therefore this splitting may not be clearly resolved experimentally. Additionally, the splitting is well below the typical uncertainty of $25\text{-}50\text{ cm}^{-1}$ for DFT phonon calculations. The model also accurately reproduces the G_i and G_b wavenumber shift of $\sim 4\text{ cm}^{-1}$ per stage, with the exception of stage 2, which shows a more significant reduction from stage 3 of 8 cm^{-1} . Compared to experiment, the model consistently underestimate the vibrational frequency of each mode by $5\text{-}60\text{ cm}^{-1}$, which is slightly beyond the typical uncertainty range of DFT phonon spectra. The size of this error is larger for lower numeral stages and may therefore be related to the overestimated wavenumber downshift, which is attributed to overestimation of the occupation of anti-bonding orbitals. This may suggest that the model features too much charge transfer from K^+ to the graphite layers, but it may

also be a consequence of the self-interaction problem in DFT that leads to overly-dispersive orbitals.




Fig. 5.14 Raman spectra around the G-peak for RH and DH KGICs and graphite. Experimental Raman spectra measured with a 512 nm laser are plotted over the DFT data in black dotted lines, for comparative purposes. Experimental Raman for stages 2-5 are ambiguous and are plotted here only as a guide. Computational Raman modes for RH cells are plotted in solid lines and DH cells in dashed lines. The left figure shows Raman modes without Lorentzian broadening. DH data is excluded from this figure to improve readability. The right figure shows Raman modes for both RH and DH cells with Lorentzian broadening. Spectra should only be compared qualitatively, not quantitatively as intensities are normalised to the maximum of each respective spectrum. Corresponding DFT parameters are reported in table 9.4.

5.5 Conclusion

I have computed Γ -point phonon spectra for stages 1-5 RH KGIC as well as stages 2-5 DH KGIC and pristine graphite. The phonon spectra for stages 1-5 with RH-style cells show moderate agreement with experiment though there are several shortcomings. Finite displacement phonon calculations return only the eigenvectors of the phonon modes and vibrational frequency of each mode. Although it is possible to identify Raman-active modes by computing the space group, it was not possible to compute the intensity of any Raman peaks, nor could any resonant-Raman phenomena be modelled. The phonon spectra present

in this work are therefore purely qualitative, and not quantitative. The wavenumber of the modelled phonon modes were also lower than what is observed experimentally, which is attributed to partial overpopulation of anti-bonding orbitals. The modelled spectra do, however, successfully reproduce the reduction in wavenumber with increasing K^+ content.

A broad range of DH cells were generated for stages 2-5 KGIC and although there was a 20 meV - 40 meV / atom range of total energies for each stage, the lowest energy cells were never more than 24 meV / atom higher than the RH cells. This shows that for even the most extreme DH cells, which features domains of only one K^+ , DH defects are energetically accessible under ambient conditions. Phonon spectra for the DH cells show poor correlation with the experimental spectra, regardless of the shape of the cell. However, phonon spectra have not been computed for DH cells featuring domains of more than one ion due to the computational cost. It is expected that the phonon spectra will become more similar to the RH spectra as the size of the domains increases. This is because larger domains lead to progressively larger fractions of the C-layers being flat, such as in RH cells. DH cells should be modelled with domains of various sizes but there is a linear relationship between the number of atoms required to define the cell and the number of atoms in a domain of finite size. This would provide valuable information about the dependence of the phonon spectra and the energy of the system on the size of the domains and may be used to make predictions about the size of the domains for each stage of intercalation. Unfortunately, computing phonon spectra for cells with domains containing as few as two ions each, is impractical due to the large computational cost.

The extent to which Raman spectra can be modelled for KGICs are limited by the models and methods available. There is no known *ab initio* method for modelling Raman intensity for metals as polarisation tensors cannot be computed for metallic systems. All Raman active phonon modes for a range of structures have been computed but the intensity of each of those modes is unknown. Furthermore, the size of the cell for which a phonon

calculation can be performed is constrained as phonon calculations scale as $\mathcal{O}(N^4)$ [191] with the number of atoms. Machine-learning potentials can significantly reduce the cost of phonon calculations but force predictions made by MLPs are not typically accurate enough for phonon calculations.

Stage	Cell Basis	X	Y	ΔE meV / atom	ΔZ Å	Average Gradient
5	h	5	6	25.1	0.811	0.110
5	h	3	10	42.4	1.657	0.136
5	h	2	15	30.2	1.486	0.081
5	o	5	3	43.6	0.855	0.134
5	o	3	5	24.0	1.661	0.158
4	h	4	6	20.0	0.809	0.110
4	h	3	8	14.6	0.433	0.102
4	h	2	12	24.5	1.452	0.099
4	o	4	3	19.9	0.878	0.138
4	o	3	4	13.9	0.970	0.114
4	o	6	2	17.0	0.947	0.222
4	o	2	6	29.0	1.518	0.119
3	h	3	6	20.1	0.831	0.113
3	h	2	9	28.7	1.485	0.135
3	o	3	3	20.6	0.810	0.127
2	h	3	4	16.9	0.107	0.022
2	h	2	6	14.2	0.860	0.117
2	h	1	12	59.1	1.035	0.070
2	o	6	1	11.9	1.009	0.138
2	o	3	2	33.7	0.708	0.167
2	o	2	3	14.1	0.859	0.135
2	o	1	6	57.7	1.039	0.081

Table 5.1 Summary of generated 2-layer DH cells. Cells constructed from a 2-atom hexagonal basis cell are labelled with a “h” and cells generated with a 4-atom orthorhombic basis cell are labelled with a “o”. The X and Y dimensions of the cell are also provided as a multiple of the basis cell dimensions, and are listed in the X and Y columns. The fifth column show the difference in energy per atom of the DH cell as compared to the RH cell of the same stage. The sixth column shows the distance between the minimum and the maximum of Z coordinate of C atoms within a single contiguous C layer. The final column shows the average gradient, which is calculated as the value of the fifth column divided by the length of the X or Y axis (whichever is longer).

Average charge (e) transfer per C atom			
Stage	All layers	Interior layers	boundary layers
1	-0.214	N/A	N/A
2	-0.083	N/A	-0.083
3	-0.056	-0.012	-0.076
4	-0.042	-0.009	-0.074
5	-0.033	-0.007	-0.073

Table 5.2 Layer-wise charge (given as a multiple of elementary charge, e) transfer in KGIC from Mulliken analysis. The first column shows the stage of intercalation. The second layer shows the charge transfer averaged over all layers, the third column shows the charge transfer averaged over all interior layers, and the fourth column shows the charge transfer averaged over all boundary layers. All layers in stage 1 are identical and are represented fully in the “All layers” column.

Chapter 6

Niobate Cathodes

6.1 Introduction

The initial success of LiCoO_2 in the early 1990's sparked interest in the search for other metal-oxide rocksalt-type structures, namely Ni^{3+} , Mn^{4+} , and Ru^{4+} . However, Nb^{5+} was frequently overlooked. Li_3NbO_4 , has been known and studied since at least 1963 [192, 193]. If Li^+ could be fully reversibly extracted from Li_3NbO_4 , the material would have a charge capacity of 452 mAhg^{-1} , but it is an insulator and consequently, it has a charge capacity of less than 30 mAhg^{-1} is observed experimentally. However, in 2015, Yabuuchi *et al* synthesised a Mn-substituted sample of Li_3NbO_4 by mechanically milling with LiMn_2O_2 . This new material, $\text{Li}_{1.3}\text{Nb}_{0.3}\text{Mn}_{0.4}\text{O}_2$, was demonstrated to have a reversible charge capacity of 300 mAhg^{-1} at 50°C . $\text{Li}_{1.3}\text{Nb}_{0.3}\text{Mn}_{0.4}\text{O}_2$ achieved its own theoretical charge capacity of 380 mAhg^{-1} on the initial discharge but suffered capacity fading due to decomposition [194]. The structure of this compound is considerably different to LiCoO_2 , however. While LiCoO_2 features Li^+ intercalated between rocksalt-like CoO_2 layers, in Li_3NbO_4 , the Li^+ ions are incorporated fully into the rocksalt-structure (see Figure 6.1). Li^+ migration is expected to be slow for these types of structures, however, Li^+ is able to percolate through the network, resulting in a discharge capacity of 200 mAhg^{-1} at 800 mA g^{-1} for Li_3NbO_4 [194].

In 2004, McLaren *et al* demonstrated improved Li^+ kinetics in substituted Li_3NbO_4 materials, as demonstrated by the increase in conductivity from $4 \times 10^{-7} \Omega^{-1}\text{cm}^{-1}$ to $5 \times 10^{-4} \Omega^{-1}\text{cm}^{-1}$ [195]. Substitution in Li_3NbO_4 have consequently been a focus of subsequent research into improving and understanding the material's kinetics [196–198].




Fig. 6.1 LiCoO_2 (left) and cubic Li_3NbO_4 (right) cells with Co shown in blue, Nb in dark green, O in red, and Li in light green.

The Na analogue, Na_3NbO_4 , has seen less research than Li_3NbO_4 [199–201]. This can be attributed to the lesser interest in Na-ion batteries (NIBs) over Li-ion batteries (LIBs), however, Na_3NbO_4 has some clear shortcomings. Na_3NbO_4 has a lower theoretical charge capacity than Li_3NbO_4 , at only 356 mAhg^{-1} , but Na_3NbO_4 is also an insulator and consequently also shows significantly reduced capacity in practice. Na_3NbO_4 shares the same rocksalt-type structure as Li_3NbO_4 , but Na^+ has an ionic radius of 1.9 \AA , which is greater than the 1.7 \AA ionic radius of Li^+ , and so slower ionic transport is expected in Na_3NbO_4 . In a similar approach to Yabuuchi *et al*, in 2017 Sato *et al* attempted to substitute the otherwise “electrochemically inactive” Na_3NbO_4 with Mn. Mechanical milling of Na_3NbO_4 of NaMnO_2 yielded $\text{Na}_{1.3}\text{Nb}_{0.3}\text{Mn}_{0.4}\text{O}_2$, an identical stoichiometry for that achieved for the Li analogue, and achieved a reversible charge capacity of 200 mAhg^{-1} [202, 203]. Yabuuchi

et al had also attempted to synthesise the Na analogue of Na_3NbO_4 by the same method, but instead observed phase separation into Na_3NbO_4 and NaMn_2O_2 [196].

The ionic conductivity of a material is intrinsically linked to its structure, while its theoretical gravimetric charge capacity is determined by its stoichiometry (assuming all charge carriers are labile). Therefore, it may be possible to synthesise a polymorph of Na_3NbO_4 with similar charge capacity but higher ionic conductivity. Prof. Kent Griffith has synthesised a number of niobates that may have battery applications [204], including a polymorph of Na_3NbO_4 in which two thirds of the material's total Na^+ are contained in Na^+ layers. Layered structures display excellent kinetics, suggesting that Na_3NbO_4 could be a strong candidate for a NIB cathode. Prof. Griffith performed ion exchange to substitute Na^+ for Li^+ to synthesise layered Li_3NbO_4 , which likewise is a strong candidate for a LIB cathode. However, ion exchange is a reversible process and the equilibrium ratio of Li^+ to Na^+ is not known directly. Solid-state NMR for the isotopes ^{23}Na and $^6,^7\text{Li}$ has been performed on the material before and after ion exchange. ^{23}Na activity becomes negligible after ion exchange, suggesting that the process has run to near-completion. This research was conducted as part of a collaboration with Prof. Griffith to verify that Na_3NbO_4 and $\text{Li}_{1-z}\text{Na}_z\text{NbO}_4$ take the expected structures by modelling NMR spectra for cubic Li_3NbO_4 and Na_3NbO_4 , as well as for layered Li_3NbO_4 and Na_3NbO_4 . See section 6.2 for further details.

Many Li and Na niobates have been reported in the literature, but not all have received equal attention. $\text{NaNb}_7\text{O}_{18}$ was first characterised in 1984 by B.-O. Marinder and M. Sundberg and has received relatively little research since [205]. $\text{NaNb}_7\text{O}_{18}$ is unfortunately unlikely to see any application as a NIB electrode due to its extraordinarily low theoretical charge capacity of $<28 \text{ mAhg}^{-1}$, but it has a potential application as a Li^+ host for LIBs. $\text{NaNb}_7\text{O}_{18}$ may be able to accommodate as many as 7 Li^+ charge carriers per formula unit, giving it a theoretical charge capacity of 185.5 mAhg^{-1} . $\text{NaNb}_7\text{O}_{18}$ has been characterised

with an *Immm* symmetry and a unit cell containing four $\text{NaNb}_7\text{O}_{18}$ formula units. The cell parameters of the unit cell are known precisely, as are the positions of the Nb and O atoms. However, while the positions of two of the Na^+ ions are known exactly, the positions of the remaining two Na^+ ions are not. They are believed to occupy two of eight possible interstitial sites (the structure and possible sites are shown in Figure 6.2). Although Prof. Griffith has demonstrated that Li^+ can transport through the material, Na^+ does not and therefore it effectively acts as an obstacle to ion transport. Accurate structural characterisation of the material will aid in understanding the ionic conductivity and consequently the performance of $\text{NaNb}_7\text{O}_{18}$ as a battery material. A similar approach to Na_3NbO_4 has been taken for structure determination. Prof. Griffith has performed ^{23}Na solid-state NMR on the material and, as part of this thesis, NMR spectra have been modelled on several potential polymorphs such that they can be compared to experiment to attempt to identify the position of the remaining two Na^+ ions.

6.2 X_3NbO_4

NMR calculations were performed for Li_3NbO_4 and Na_3NbO_4 with both a layered and cubic structures (visualised in Figures 6.3 and 6.4, respectively). The chemical shieldings for Li_3NbO_4 and Na_3NbO_4 are presented in Figures 6.5 and 6.6. The cubic cells contained only a single chemical environment and so only a single peak is observed in the chemical shielding spectra for both Li_3NbO_4 and Na_3NbO_4 . The spectra for the layered structures were more complex and each featured six unique chemical environments. Each spectrum contained one peak with relative intensity 1, three peaks with relative intensity 1/2 and two peaks with relative intensity 1/4. The relative shielding of the peaks in the spectra was not consistent between Li and Na, with the high intensity peak appearing with the fourth highest chemical shielding in the Li spectra, but with only the second highest chemical shielding in the Na spectra.




Fig. 6.2 Model cell showing expected structure of $\text{NaNb}_7\text{O}_{18}$ unit cell including additional Na^+ ions in potential interstitial sites. Potentially occupied octahedral interstitial sites are labelled a, b, c, and d when adjacent to a cell edge parallel to the X-axis and α , β , γ , δ with adjacent to the cell edge parallel to the Y-axis. The remaining unlabelled Na^+ ions occupy known square planar sites. Nb atoms are shown in green, O in red, and Na in yellow.




Fig. 6.3 Li_3NbO_4 unit cells for cubic (left) and layered (right) structures. Nb atoms shown in dark green, O in red, and Li in light green.




Fig. 6.4 Na_3NbO_4 unit cells for cubic (left) and layered (right) structures. Nb atoms shown in dark green, O in red, and Na in yellow.




Fig. 6.5 Li_3NbO_4 isotropic chemical shieldings for layered (top) and cubic (bottom) structures. The chemical shielding spectra for layered Na_3NbO_4 contains six peaks with the intensity ratio 2:1:2:2:4:1. The chemical shielding spectra for layered Na_3NbO_4 contains one peak. Li_3NbO_4 cells are overlaid on their respective spectra with Nb shown in dark green, O in red, and Li in light green.

^{23}Na has a spin of $3/2$ and hence ^{23}Na NMR spectra show quadrupolar coupling. Consequently, the chemical shieldings presented in Figure 6.6 are not directly comparable with experimental ^{23}Na spectra. The quadrupolar coupling for each ^{23}Na environment has been computed with CASTEP and then Prof. Griffith has applied this coupling to the modelled spectra. Furthermore, the spectra reported in Figures 6.5 and 6.6 return the chemical shielding on the X-axis, while experimental spectra show the chemical shift. NMR spectrometers measure resonant frequencies which are dependent on the strength of the magnetic field in the spectrometer. For consistency, NMR spectra are usually expressed in terms of the chemical shift with respect to the resonant frequency of some reference compound, which is independent of the spectrometer. In this instance, Prof. Griffith has applied a shift by referencing the modelled spectra for the two Na_3NbO_4 compounds and the $\text{NaNb}_7\text{O}_{18}$ candidate structures (see section 6.3) to his experimental ^{23}Na .




Fig. 6.6 Na_3NbO_4 isotropic chemical shieldings for layered (top) and cubic (bottom) structures. The chemical shielding spectra for layered Na_3NbO_4 contains six peaks with the intensity ratio 1:2:2:4:1:2. The chemical shielding spectra for layered Na_3NbO_4 contains one peak. Na_3NbO_4 cells are overlaid on their respective spectra with Nb shown in dark green, O in red, and Na in yellow.

As shown in Figure 6.7, the modelled ^{23}Na spectrum for layered Na_3NbO_4 is highly consistent with the experimental ^{23}Na collected by Prof. Griffith. Every feature of the of the experimental spectrum can be clearly assigned to individual Na environments, as shown by the deconvolution of the modelled spectra in the same figure. It therefore seems highly likely that Prof. Griffith succeeded in layered Na_3NbO_4 .

6.3 $\text{NaNb}_7\text{O}_{18}$

6.3.1 Structure Searching Method

A cell of $\text{NaNb}_7\text{O}_{18}$ may be defined with no fewer than four formula unit, such as in Figure 6.2. This unit cell is believed to contain two Na^+ ions in square planar coordinations, which in this representation appear at the centre and corners of the cell. There are two additional Na^+ ions which are believed to occupy two of eight interstitial sites of a octahedral geometry. The octahedral sites form channels in the Z-direction that may facilitate Li^+ transport in Li-doped $\text{NaNb}_7\text{O}_{18}$. In order to determine the correct positions of the Na^+ ions in $\text{NaNb}_7\text{O}_{18}$, all possible structures have been generated and optimised. Firstly, the constraint is imposed that the models shall include no more than one unit cell. Such a constraint does not allow for modelling a structure in which Na^+ ions are irregularly distributed in the Z-axis as each unit cell contains only one interstitial site per channel. This constraint is necessitated by the computational cost associated with optimising the cells and modelling the NMR spectra due to the large numbers of ions in each unit cell, and due to the relatively large number of k-points. The number of unique cells increases exponentially with the size of the supercell, and the cost of the DFT also increases with the size of the cell. Structure searching therefore focused on a single unit cell, in which the Na^+ ions are regularly distributed, leaving six of the eight channels entirely empty. Six additional calculations were later performed on




Fig. 6.7 Figure to compare experimental ^{23}Na NMR spectrum of layered Na_3NbO_4 with modelled Na NMR spectrum. The top spectrum is experimental, the bottom six spectra show the modelled spectra deconvoluted by unique Na environment, and the fit shows the sum of these spectra. Prof. Griffith has applied a shift to the chemical shieldings by referencing the modelled spectra to his experimental spectra. The modelled Na NMR spectrum of layered Na_3NbO_4 is highly consistent with the experimental spectrum of Na_3NbO_4 , showing that Prof. Griffith was successful in synthesising layered Na_3NbO_4 . Figure provided by Prof. Griffith and reproduced with permission. Specific chemical environments are shown in figure 9.4.

$1 \times 1 \times 2$ supercells so that the impact of disorder can be assessed. This is described in greater detail later in the chapter.

Due to the highly constrained phase space that is explored in this research, an AIRSS approach is excessive. If the Na^+ in the octahedral sites are ignored (because they are unknown), the unit cell has *Immm* symmetry, which includes numerous symmetry features such as an XY, XZ, and YZ mirror planes and 2-fold rotational symmetry around the Z-axis. Consequently, for a single unit cell, there are only ten possible unique cells. Therefore, a structure *enumeration* approach is taken rather than a structure *searching* one. Due to the limited number of unique structures, a descriptive nomenclature is defined based on the positions of the Na^+ ions such that each structure has an unambiguous descriptive name. For the unit cell in Figure 6.2, each octahedral site that may contain a Na^+ ion has been labelled with a Latin or Greek letter, with Latin letters being used for sites along the X-edge of the cell and Greek letters for sites along the Y-edge of the cell. Within the XY plane, these sites may be thought to form a square-shaped “network” of four sites through coordination to shared O atoms. The unit cell defined in Figure 6.2 contains two of these networks. One network contains only Greek-labelled Na^+ sites and the other only Latin-labelled sites. Any unique arrangement of Na^+ ions may have multiple valid names under this scheme, depending on the viewer’s frame of reference. For example, “ab”, “cd”, “ $\alpha\beta$ ”, and “ $\gamma\delta$ ” are equivalent. “a δ ”, “ αd ”, “b γ ”, and “ βc ” are also equivalent to each other. A cell for each Na^+ arrangement was constructed manually and then geometry optimised.

$\text{NaNb}_7\text{O}_{18}$ cells generated during this work that comprised of four formula units (4FU) were less than 4 Å in length in the Z-axis but over 14 and 26 Å long in the X and Y axes. Atoms consequently experienced a significantly stronger interaction with their periodic image in the Z-direction than in the other axes. To assess the impact of distributing Na^+ ions across multiple channels, cells were generated which were expanded in the Z-axis to form eight formula units (8FU) cells. The energy per atom of $1 \times 1 \times 2$ supercells is not expected to

differ from their unit cells (within a typical degree of uncertainty from a DFT calculation with a random seed), but energy per atom of the 8FU cells created from two different 4FU cells provides insight into the effect of Na⁺-Na⁺ interaction in the Z-direction. Due to large computational cost of performing geometry optimisations and NMR calculations on 8 FU cells, only two were tested.

To estimate the uncertainty of the energy prediction, the difference in energy between cubic and layered Na₃NbO₄ has been plotted with respect to k-point spacing and planewave cutoff. Figure 6.8 shows that at the chosen k-point spacing, 0.03 Å⁻¹, the energy difference differs by no greater than 0.01 meV / atom for denser k-point grids, and for the chosen planewave cutoff, 800 eV, the energy difference differs by no greater than 0.02 meV / atom for higher energy cutoffs. The convergence for three different SCF random seeds has been plotted, which reveals its own uncertainty of no greater than 0.01 meV / atom. The total uncertainty from the k-point spacing, planewave, and SCF random seed is therefore estimated to be less than 0.03 meV / atom.




Fig. 6.8 Convergence of energy difference per atom between layered and cubic Na₃NbO₄ with respect to k-point spacing and planewave cutoff.

6.3.2 Structure Analysis

The lowest energy cell was $a\delta$ and the highest energy polymorph was $\alpha\beta$, however these differed in energy by only 3.28 meV / atom, suggesting that all polymorphs may be accessible at room temperature. The four lowest energy cells ($a\delta$, $\alpha\gamma$, $\alpha\alpha$, and $\alpha\beta$) are all within 0.12 meV / atom of $a\delta$, as seen at the bottom of Figure 6.9. All these cells feature one Na^+ ion in the top left site (with respect to Figure 6.9), and the second in one of the four sites on the left or right. This arrangement minimises the number of Na^+ ions in the second and third Na^+ neighbour shell because the two Na^+ ions are in different four-site networks. The number of ions in the first Na^+ neighbour shell is fixed across all polymorphs as the first shell shows the distance between an atom and itself in a periodic image in the Z-direction. For practical purposes, these four low energy cells may be treated as degenerate as they are energetically accessible within less than 2 K of one another. Further, each cell is within two degrees of uncertainty of one or more other cells.

The next three cells differ more significantly in energy. $\alpha\delta$, $\alpha\gamma$, and $\alpha\beta$ are 1.92, 2.27 and 3.28 meV / atom higher in energy than $a\delta$. These higher energy polymorphs contain similar numbers of ions in the second and third Na^+ neighbour shells, but the relatively high energy of $\alpha\beta$ is attributed to the continuous chain of Na ions along the X-axis. A complete summary of the seven candidate unit cells with their relative energy can be found in Figure 6.10.

The energies of two 8FU unit cells composed of two 4FU unit cells has been compared (Figure 6.11). The first cell, labelled 8FU-A, comprises $a\delta$ and $\alpha\delta$, which are the same structures but rotated 180°. The second cell, labelled 8FU-B, comprises $\alpha\beta$ (equivalent to $\alpha\gamma$) and $\alpha\gamma$. 8FU-A shows 2_1 screw axis symmetry in the Z-direction and so if Na^+ ions did not experience each other in the Z-direction, it is expected that 8FU-A would have the same energy per atom as $a\delta$. Instead, there is an increase in energy per atom of 0.68 meV / atom. In contrast, 8FU-B shows no symmetry and is comprised of two completely different unit cells. Again in the hypothetical situation where Na ions do not experience each other in the




Fig. 6.9 Na chemical shieldings for the seven possible polymorphs of the NaNb₇O₁₈ unit cell. Polymorphs are ordered according to their energy as compared to the lowest energy polymorph. They are also coloured proportionally to this difference. Polymorphs are labelled according to the nomenclature set out in Figure 6.2. Each peak is labelled according to the Na ion associated with it, where '1' and '2' represent the a Na ion in a square planar site in the corner and central positions in figure 6.2, and the other Na sites are labelled according to the same figure.




Fig. 6.10 Candidate unit cells for $\text{NaNb}_7\text{O}_{18}$, labelled according to the scheme in Figure 6.2. The difference in energy between the cell and the lowest energy cell is given below each image. Nb atoms are shown in green, O in red, and Na in yellow.

Z-axis, it is expected that the energy per atom of 8FU-B would be equal to the mean energy per atom of ab and $\alpha\gamma$. Instead, it is observed that the energy per atom is lower than either constituent cell, and only 0.35 meV / atom higher than the lowest energy cell, $\alpha\delta$. It therefore seems that homogeneous occupation of Z-axis channels, as in the case of 4FU unit cells, has neither a consistently positive nor consistently negative impact on the energy of the cell, but rather that the energy of the cell has a complex relationship with the distribution of the Na ions in 3-dimensions.

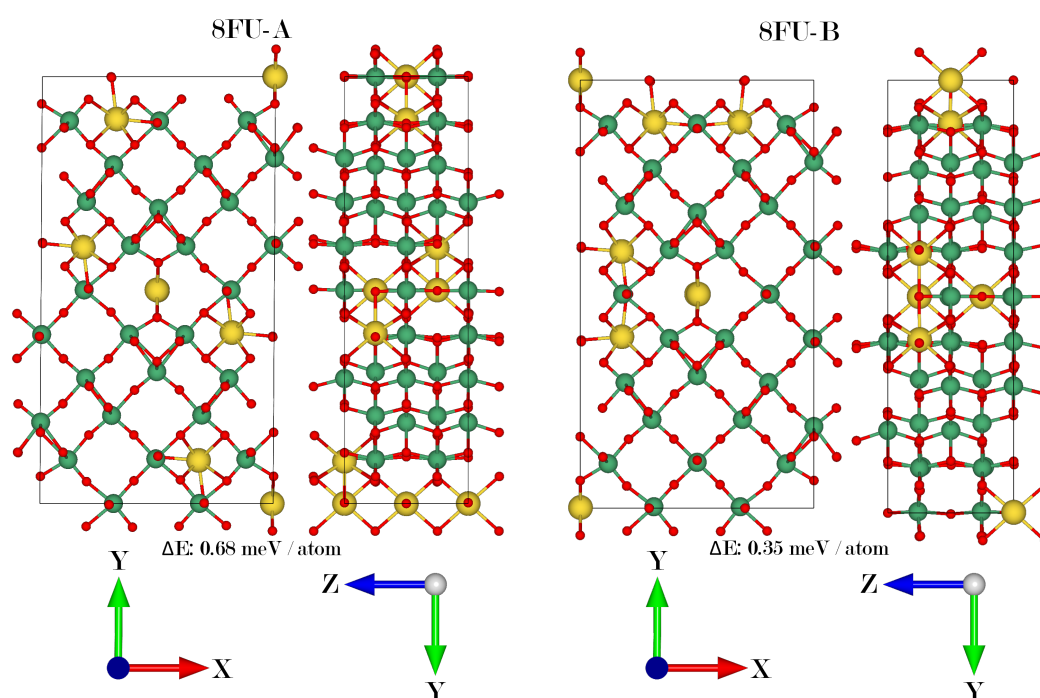


Fig. 6.11 Visualisation of 8FU-A (left) and 8FU-B (right) cells from the Z and X axes. Each cell is higher in energy than the $\alpha\delta$ by 0.68 and 0.35 meV / atom, respectively. 8FU-A is higher in energy than its constituent cells, while 8FU-B is lower in energy than its constituent cells.

6.3.3 NMR

The chemical shieldings of the seven polymorphs are unlikely to allow for their identification through experimental NMR as the peaks shift by less than 1 ppm between polymorphs. This

is somewhat less than the expected sensitivity of any experimental measurements. Only one of the polymorphs has a spectrum that differs significantly from the others and this is $\alpha\beta$. Here, the spectrum only contains three unique environments, resulting in two peaks of equal intensity and a third peak with double the intensity of the other two (twice as many atoms are in that environment). However, the two low ppm peaks for the other six polymorphs show a splitting on the order of 0.1 ppm, and are therefore likely to appear as a single peak experimentally.

The similarities of the chemical shieldings also poses a problem for the accuracy of this DFT model. For the three high energy polymorphs, additional structure files were generated which contained Na^+ ions in different sites but for which the resulting structure was equivalent to another which had already had an NMR calculation performed on it. These three polymorphs were $\alpha\beta$, $\alpha\gamma$, and $\alpha\delta$. Geometry optimisations were performed on each of these structures and modelled the NMR spectrum. While the shielding peaks are expected to differ slightly for two equivalent structures from different geometry optimisations, it is found that the difference in shielding was of a similar scale to the difference in chemical shieldings between polymorphs, approximately 1 ppm. The uncertainty of the model therefore appears to be too large to be able to differentiate the possible Na^+ arrangements. In order to ensure that self-interaction between periodic images was not significantly impacting the error in the models, $1 \times 1 \times 2$ supercells were created of two geometry optimised but equivalent structures which had the greatest difference in shieldings ($\alpha\delta$ and $\alpha\gamma$). NMR spectra were modelled for these supercells without performing any further geometry optimisation. Figure 6.12 shows that the supercells had slightly lower chemical shieldings but did not differ significantly from the unit cells. It is also noteworthy that the spectra of each of the supercells was more similar to the spectra of their respective unit cells than to any other spectrum, suggesting that the difference in chemical shielding was caused by differences in atomic positions, rather than a self-interaction problem.

Making an objective comparison of two structures is difficult but pair distribution functions (PDF) allow for the distribution of ions throughout the cell to be compared, highlighting any significant differences in inter-atomic distances. From visualisation of the cell, it is clear that the overall geometry of $\alpha\delta$ and ad are the same and the cells cannot be differentiated by eye. Any difference between the cells therefore must be small changes in the atomic coordinates, which is expected to be reflected in a PDF. However, Figure 6.13 shows that the PDFs of the two structures are seemingly identical below 8 Å and even beyond this point, differences are negligible. The two structures also only differ in energy by 0.09 meV / atom, suggesting that these NMR spectra are extremely sensitive to geometry. It may be possible to reduce the uncertainty by reducing the k-point spacing and the geometry optimisation tolerances, but this is expected to be prohibitively computationally expensive.

As described in section 6.2, Prof. Griffith has applied quadrupolar coupling and to the modelled spectra and applied a linear shift to the chemical shielding by referencing between the modelled spectra and experiment. As shown in Figure 6.14, the modelled ²³Na spectrum for NaNb₇O₁₈ is generally consistent with the experimental ²³Na collected by Prof. Griffith. The experimental spectra may be described in terms of four peaks at -8 ppm, -25 ppm, -38 ppm, and -60 ppm. The high intensity peak at -8 ppm was consistent with the signal from the ²³Na⁺ ions in octahedral sites, but full-width-half-maximum of those peak was approximately 5 ppm, which was more than twice the difference in chemical shift between the different ²³Na configurations in the NaNb₇O₁₈ candidate cells. These different environments therefore cannot be resolved experimentally. The peaks at -38 ppm, and -60 ppm lined up well with the peaks from the ²³Na ions in square-planar sites, and the experimental peak at -25 ppm has been attributed to a NaNb₃O₈ impurity. Given the constraints of the model and the resolution of the experimental measurements, there is deemed to be good consistency between the model cell and the cell synthesised by Prof. Griffith.




Fig. 6.12 Na chemical shieldings for $\alpha\delta$ and ad, as well as their supercells. Where the unit cell is labelled with $1 \times 1 \times 1$ and the supercell with $1 \times 1 \times 2$. All cells have the same general geometry but show distinct chemical shieldings. The chemical shieldings of the supercells are more similar to the chemical shieldings of their respective unit cells than to any other cell.




Fig. 6.13 Na pair distribution function (PDF) for $\alpha\delta$ and ad post-geometry optimisation. The PDFs appear identical below 8 Å and are very similar above this point.

6.4 Conclusion

As part of a collaboration with Prof. Griffith, the Na NMR spectra of two Na-niobate compounds which were experimentally synthesised by Prof. Griffith have been modelled. In the first instance, Prof. Griffith attempted to synthesise a previously unseen layered polymorph of Na_3NbO_4 through ion exchange with the layered polymorph of Li_3NbO_4 . Using layered Li_3NbO_4 as a reference for the structure, a geometry optimisation was performed on the structure, and the chemical shieldings and quadrupolar couplings were modelled for each ^{23}Na environment in the cell. The modelled ^{23}Na NMR spectrum for Na_3NbO_4 was highly consistent with the experimental spectrum, suggesting that Prof. Griffith succeeded in synthesising layered Na_3NbO_4 .

Prof. Griffith has also synthesised $\text{NaNb}_7\text{O}_{18}$ and has experimentally identified the positions of the Nb and O atoms in the unit cell, which contains four formula units. Two Na ions were believed to occupy square planar sites and the remaining two ions were in unknown positions. Cells for $\text{NaNb}_7\text{O}_{18}$ containing all possible unique arrangements of Na ions in octahedral sites for a cell containing four formula units (4FU) have been generated. For each of these cells, the Na NMR spectra were modelled such that they can be compared to experimental spectra in order to determine how Na ions are distributed through the interstitial sites. The difference in chemical shielding between spectra was less than 2. This value was smaller than the uncertainty within the model and also too small to be resolved experimentally with ^{23}Na NMR; therefore it was not possible to unambiguously assign any one of the generated cells to the experimental spectra. It is clear from this data that the sensitivity of Na^{23} NMR is insufficient to resolve the occupied interstitial site. Insight to the structure of $\text{NaNb}_7\text{O}_{18}$ may be gleaned from the relative energy of the cells generated during this work. Of the seven 4FU unit cells, the highest energy cell was only 3.28 meV / atom higher in energy than the lowest energy cell, and three of the cells were within 0.12 meV / atom of the lowest energy cell. Two eight formula unit cells (cells that were not supercells

of any 4FU unit cells) were also generated and these were shown to have an energy 0.68 and 0.35 meV / atom higher in energy than the lowest energy cell. In one case, the cell was higher in energy than its constituent cells and in the other case, it was lower in energy than its constituent cells, suggesting that homogeneous occupation of the Z-channels has no consistent impact on the energy of the cell. Given the small energy range (3.28 meV / atom) for all $\text{NaNb}_7\text{O}_{18}$ cells generated as part of this work, all cells are expected to be energetically accessible at room temperature. It is therefore expected that Na to be disorderly distributed throughout the structure. The partial occupancy of all interstitial sites by Na necessitates Li^+ transport by percolation, a slower process than direct migration through the lattice. Furthermore, although the modelled NMR spectra were generally consistent with the experimental spectra, the modelled spectra did not reproduce the peak at -25 ppm, suggesting inconsistency between the model cell and the synthesised cell. Given that the octahedral sites had been thoroughly explored during this work, it should now be considered whether the square planar sites both accommodate Na^+ ions as previously believed. This work will be expanded by modelling $\text{NaNb}_7\text{O}_{18}$ cells which contain two Na ions offset from the square planar sites into the neighbouring 11-coordinate sites.




Fig. 6.14 Figure to compare experimental ^{23}Na NMR spectrum of layered $\text{NaNb}_7\text{O}_{18}$ with modelled Na NMR spectrum. The top spectrum is experimental, the bottom two spectra shows the modelled spectra deconvoluted by the two unique Na environments in square-planar (which are the unlabelled Na ions in figure 6.2) and octahedral sites (which are the Na ions labelled a , b , c , d , α , β , γ , and δ in figure 6.2), and the fit shows the sum of these spectra. Each of the octahedral sites considered in this work show highly similar chemical shieldings and quadrupolar coupling and hence they cannot be resolved. The octahedral environments are not reported separately for this reason. Prof. Griffith has applied a shift to the chemical shieldings by referencing the modelled spectra to his experimental spectra. Experimental peaks at -8 ppm, -38 ppm, and -60 ppm are consistent with the three peaks in the modelled spectrum. The experimental spectrum also contains a peak at -25ppm, which is attributed to a NaNb_3O_8 impurity. This suggests that the proposed structure of $\text{NaNb}_7\text{O}_{18}$ is a good fit with the material synthesised by Prof. Griffith. Figure provided by Prof. Griffith and reproduced with permission.

Chapter 7

Conclusion

In this thesis, high-throughput structure searching of functional materials has been approached in a broad range of scenarios. This included searching for a stable ternary cathode material for a system where no stable ternary compounds were known, searching for point defects in a low dimensional SPE where previously only a limited range of defects were known, comparing two competing structural models for a layered anode material and comparing their modelled Raman spectra to experimental spectra, and modelling NMR spectra to verify the structure of two experimentally synthesised compounds for which the structure was not fully understood. In two of these projects, random structure searching was applied to explore a wide phase space, while the other two projects required a more targeted approach. In all projects, density-functional theory (DFT) has been used, which allowed for accurate modelling of energies, electronic structures, phonon spectra, and NMR spectra, and in two of the projects, machine learning (ML) models were used to accelerate geometry optimisations by three orders of magnitude. Each project posed its own challenges and required separate, but related, approaches.

In Chapter 3 *ab initio* random structure searching (AIRSS) and machine learning random structure searching (MLRSS) were used to to explore Li-Ni-S ternary phase space in search of a new high charge capacity cathode for Li-ion batteries. In the first instance,

AIRSS [27] was used, which involved randomly generating structures and then performing a geometry optimisation using DFT for the energy, force, and stress evaluations. The most computationally demanding step of AIRSS is the geometry optimisation due to the large computational cost of performing DFT calculations at every step. In MLRSS, the structure searching process was significantly accelerated by replacing DFT with a machine learning potential (MLP), which can perform energy, force, and stress evaluations three orders of magnitude faster. The MLP was acquired through an iterative process of short MLRSS runs, followed by DFT singlepoint calculations on the optimised structures, adding the DFT data to an existing database, and training the next generation MLP with the PaiNN architecture [76]. The AIRSS and MLRSS data set was further expanded by collating structures that contained similar elements from the Materials Project [117], the Open Quantum Materials Database [118], and a recent publication by Cen *et al* [95], and then swapping these elements for Li, Ni, and S. The Li-Ni-S data set contained over 150,000 structures but no ternary compounds were identified on the convex hull. The structures which were identified on the convex hull were generally consistent with what has been reported in the literature, which are the three constituent elements, and the Li_2S , Ni_3S_2 , NiS , Ni_3S_4 , and NiS_2 binaries. NiS conversion-type cathodes are widely studied in the literature and the experimentally-observed conversion process is consistent with the predictions made by the phase diagram. In addition, the experimental voltage curve for lithiation of NiS is consistent with the modelled voltage curve. NiS_2 is also a conversion-type cathode but with a gravimetric charge capacity of 712 mAh g^{-1} against the charge capacity of 512 mAh g^{-1} for NiS . It has been less widely studied, however, due to its thermal instability. The modelled phase diagram finds NiS_2 on the convex hull, and that when lithiating, conversion into Ni_3S_4 and Li_2S is expected, but also that lithiated Ni_3S_4 should convert to NiS and Li_2S . Furthermore, the modelled voltage curve overestimates the experimental voltage curve for this system by 0.2 V, but is within acceptable tolerances for a model such as this. Xu *et al* suggested that lithiated NiS_2

firstly forms an unknown Li_xNiS_2 [106] compound but no such compound was discovered during the searches. None of the identified Li_xNiS_2 compounds were closer than 59 meV / atom to the convex hull, and therefore none of the discovered compounds are likely to be stable under ambient conditions. In lieu of any stable Li_xNiS_2 , it appears that Li-Ni-S ternary compounds will not see application as battery materials. However, NiS_2 has a charge capacity of 712 mAhg^{-1} , approximately 2.5 times greater than that of LiCoO_2 . It has not yet been implemented commercially due to its poor stability but a new polymorph of NiS_2 has been identified with a formation energy that is predicted to be lower than formation energy of the polymorph that is typically used for NiS_2 cathodes. It should be determined whether this material can be synthesised, and if so, whether it is stable, how quickly its capacity fades during cycling, and how quickly it can charge and discharge.

In the following chapter AIRSS was applied but instead of new phases, searches were performed for point defects in monolayer hBN. While most single photon emitters (SPEs) require cryogenic temperatures, hBN is a room temperature SPE, potentially making it significantly cheaper and simpler to implement than other SPEs. hBN's ability to act as an SPE has been attributed to point defects but the exact defect(s) is unknown. If it can be identified, it may be possible to engineer these defects into hBN monolayers to create energy-specific and high intensity SPEs. So far, attempts to identify the defects responsible have been limited to simple substitutional and vacancy defects, but a far broader range of defects are possible. Through AIRSS, it was possible to generate over 5000 intrinsic and extrinsic defects, revealing many low energy structures that have not previously been considered in studies of hBN SPEs. To assess the stability of these defects, the defect formation energy was calculated and then normalised to the minimum number of possible defect sites for each defect. This allows for the formation energies of defect with different stoichiometries to be compared in such a way that is independent of the size of the pristine hBN around the defect, and independent of any atomic rearrangements that may lead to a change in the

number of defect sites. For computational efficiency, all AIRSS was performed in $5 \times 5 \times 1$ monolayer hBN supercells, but all intrinsic defects and the lowest energy defect of each stoichiometry were optimised again in $15 \times 15 \times 1$ to minimise strain in the cell. The lowest energy defects will now be shared with our collaborators, Prof. Bartomeu Monserrat, and Mr. Siyu Chen who will model the optical absorption and emission spectra by solving the Bethe-Salpeter equation. Where spectra show absorption/emission peaks at similar energies to the experimentally-measured energy of emitted single-photons, the spectra will then be more comprehensively compared to experimental spectra with the hope of identifying the point defect that gives rise to the experimentally-observed SPE phenomenon. In addition to this, it may be possible to identify point defects that give rise to desirable SPE phenomena that has not been observed experimentally. For any of these point defects with desirable optical properties, it will next need to be determined how they can be engineered and in what concentration, as this influences the intensity of the single-photon emission.

In Chapter 5, Raman spectra for stage 1-5 KGIC were modelled within both the Rüdorff-Hoffman (RH) and Daumas-Hérold (DH) models for graphite-intercalation compounds (GIC). Potassium-GIC (KGIC) cells with the RH models are rigorously defined, allowing for a cell for each stage to be manually constructed and then geometry optimised to the global minimum. However, the value of several variables which define DH cells, including domain sizes and the curvature of graphite layers, were unknown and therefore a structure searching method was necessary. To this end, a bespoke algorithm was applied which automated the construction of candidate KGIC cells through a Python3 script and then optimised the structures with an MLP. This approach allowed for the generation of DH cells with size and complexity that is inaccessible through DFT geometry optimisations due to the large computational demands. However, the errors on the force predictions from the MLP were too large to model phonon spectra to a sufficient level of accuracy; meaning that DFT was required for performing phonon calculations. Consequently, the DH cells were constrained

to two layers and a domain size of one ion so that phonon spectra could be modelled in a practical time frame and with a practical amount of computing resources. The DH cells were each higher in energy than the RH cells of the same stage, but by less than 24 meV / atom, suggesting that they are energetically accessible at room temperature. Furthermore, the difference in energy between RH and DH cells is largely attributed to strain in the graphite layers caused by the forced curvature of those layers. For DH cells with larger domain sizes, strain is reduced and therefore the DH cells considered in this work act as an upper bound for the energy difference between pristine RH cells and RH cells that naturally form with DH defects. The DH cells with the lowest formation energy were generally those in which the length of the X-axis was similar to the length of the Y-axis, and these cells were subsequently used for modelling phonon spectra. All RH and DH cells were low-dimensional metals due to charge transfer from K pushing the Fermi level above the graphene Dirac cone and into the conduction band. *Ab initio* modelling of Raman intensities for metals is not currently possible and so a phonon spectrum for each cell has been computed and then used the space group of each phonon eigenvector to identify Raman-active modes given that Γ -point phonon modes which induce a change in electric-dipole polarisability are Raman-active. The Raman spectra of the RH cells showed good agreement with experimental spectra provided by our collaborators, Prof. Mauro Pasta and Dr. Isaac Capone, but the vibrational frequencies were underestimated, which is attributed to overpopulation of the anti-bonding HOMO. Additionally, there are a number of factors which contribute to the error in the phonon spectra which arise from the inexact implementation of DFT. In particular, phonon spectra are intrinsically linked to the geometry of a structure and the forces experienced by its atoms during vibration. Since this implementation of DFT does not give exact energy evaluations (see section 2.2 for details), there are likely to be errors in the shape of the potential energy surface and hence the forces. Since geometry optimisations are performed using these same energies and forces, the geometry is also likely to be partially inaccurate. [206] The DH cells, in

all cases, show poor agreement with experiment, suggesting that, if present, the domains are very large and therefore DH-type cells may be viewed as defected RH cells, rather than a model for an independent phase. The impact of DH defects on KGIC Raman spectra may be better understood by modelling phonon spectra with KGIC cells with larger domain sizes, but for this to be practically possible it requires the development of an MLP with more accurate force predictions, or the development of a new ML-approach to phonon/Raman modelling. Furthermore, if an MLP can be trained that can reliably predict the energy of a large DH cell, more insight can be gained as to the relationship between stage numeral and domain size.

Lastly, two niobate compounds were characterised with NMR spectroscopy as part of a collaboration with Prof. Kent Griffith. Prof. Griffith has synthesised $\text{NaNb}_7\text{O}_{18}$ and a layered polymorph of Na_3NbO_4 which shares a structure with Li_3NbO_4 . Cubic Na_3NbO_4 has previously been experimentally observed and considered for battery applications but the material's structure is not conducive to good ion transport kinetics, which leads to slow discharge. A layered polymorph is likely to avoid this problem but until now, layered Na_3NbO_4 had not been observed experimentally. Prof. Griffith has synthesised layered Na_3NbO_4 by firstly synthesising layered Li_3NbO_4 and then performing ion exchange to replace Li^+ with Na^+ . To verify the structure of the new material, Prof. Griffith has measured ^{23}Na NMR spectra for layered and cubic Na_3NbO_4 . A layered Na_3NbO_4 cell was then geometry optimised, based on the structure of layered Li_3NbO_4 , and modelled the ^{23}Na NMR spectra of layered and cubic Na_3NbO_4 . The modelled ^{23}Na NMR spectrum, when broadened with quadrupolar coupling, shows excellent correlation with the experimental ^{23}Na NMR spectrum collected by Prof. Griffith, providing strong evidence that Prof. Griffith succeeded in synthesising layered Na_3NbO_4 . With a theoretical charge capacity of 356 mAhg^{-1} , Na_3NbO_4 exceeds the theoretical charge capacity of LiCoO_2 by almost 30%. However, to determine whether the material may be commercially competitive with LiCoO_2 as a Li-ion battery cathode, more must be learned about the material's properties, such as its

cycling stability and discharge speed. A similar approach has been taken for $\text{NaNb}_7\text{O}_{18}$. In this instance, the structure has been experimentally characterised but the positions of two of the four Na^+ ions in the proposed unit cell are unknown. The unit cell of $\text{NaNb}_7\text{O}_{18}$ contained 8 interstitial octahedral sites that could possibly accommodate the two unknown Na^+ ions. As part of a very restricted structure searching approach, all possible unique Na^+ configurations for a single cell were enumerated, geometry optimised each of them, and modelled their NMR spectra. The energy of these 7 unique cells differed by no more than 3.28 meV/atom, suggesting that each of these polymorphs are energetically accessible under ambient conditions. It is therefore proposed that the Na^+ ions are disorderly distributed throughout the octahedral sites. Prof. Griffith has broadened the modelled NMR spectra according to the calculated quadrupolar couplings, and the resultant spectra were compared to his experimental spectra. The difference between the spectra for the 7 different $\text{NaNb}_7\text{O}_{18}$ polymorphs was too small to be experimentally resolved, but if, as expected, the Na^+ ions are disorderly distributed, the experimental spectrum would show an average of those peaks, and therefore it is not necessary to resolve each polymorph. The experimental spectrum showed four peaks at -8, -25, -38, and -60 ppm. The peak at -8 ppm correlated well with the modelled peak associated with the Na^+ ions in octahedral sites, and the peaks at -38 and -60 ppm correlated well with the modelled peaks associated with the Na^+ ions in square-planar sites. The experimental peak at -25 ppm is entirely absent in the modelled NMR spectrum, and therefore can not yet be explained. This suggests that the model cells are inconsistent with the material synthesised by Prof. Griffith but given that the modelled spectra reproduces three of the four peaks so well, it is unclear as to where the error in the structure may be. One possible explanation is that the Na^+ ions do not occupy the square-planar sites as originally believed, but instead are positioned slightly off-site in an 11-coordinate environment. This alternative structure should be investigated by performing further geometry optimisations and modelling the NMR spectra. $\text{NaNb}_7\text{O}_{18}$ is being considered as a host structure for intercalation of

Li^+ to serve the role of a cathode in Li-ion batteries. The theoretical charge capacity of the compound depends on the ratio of Li^+ ions that can be reversibly intercalated, but the theoretical charge capacity is expected to be less than that of LiCoO_2 . If the material is to be commercially competitive, more will need to be learned about the material's properties, such as its cycling stability and discharge speed.

I have demonstrated that computational structure searching is a powerful and effective method of furthering our understanding of functional materials, both when used independent of experiment and when used alongside experiment. In addition, the success of the application of MLPs in Chapters 3 and 5 suggests that MLPs will play an important role in structure searching in the future. Although MLPs require an existing data set of *ab initio* or experimental data to train on, they can perform energy, force, and stress evaluations significantly faster than DFT, and with less demand for RAM. This can allow for modelling of much larger cells and hence more complex systems. While some structure searching methods are more flexible than others, there is no approach that is appropriate for every problem. Effective structure searching requires careful consideration of the range of structures that should be considered and the most efficient approach to generate and optimise them. Over the lifetime of computational atomistic modelling, the structure searching methods available to scientists has expanded from simple quasi-newton geometry optimisations to a plethora of algorithms inspired by physical phenomena, life, and statistics, and have been bettered by our improved understanding of phase space. As new structure searching methods and faster atomic models are developed, and as computing power grows, computational structure prediction will become an increasingly powerful tool in science.

Chapter 8

Further Work

Further work in the search for Li-Ni-S cathodes would benefit from the input of an experimental collaborator. If at all possible, it should be verified whether a stable Li_xNiS_2 compound does indeed exist by attempting to reproduce it in a laboratory environment. If it can be identified, any specific details about its geometry or specific stoichiometry would greatly aid in further structure searching by narrowing the phase and stoichiometric space that needs to be explored. If Li_xNiS_2 cannot be reproduced, no further structure searching should be conducted along that tie line as it has already been thoroughly explored. Additionally, further data should be collected concerning the stability of the newly discovered NiS and NiS_2 polymorphs. A finite temperature convex hull may be modelled by calculating the Gibbs free energy with the inclusion of harmonic vibrational contributions for competing polymorphs, which would provide some insight as to the relative stability of these polymorphs at a given temperature. [119] If the material can be synthesised in the laboratory and demonstrated to allow for reversible reduction by Li^+ , the material may have potential for implementation in LIBs.

The Raman spectra modelled as part of this work show that DH defects are likely to represent only a small fraction of the overall structure of KGICs. However, no claim is made as to what this fraction is. DH cells are largely impractical to model with DFT for

domain sizes greater than one and for more than two layers, but it has been demonstrated that geometry optimisations can be performed on significantly more complex cells through the use of MLPs (see Figure 5.8). The MLPs used in this work were trained primarily on forces and the accuracy of their energy predictions were not verified. If an MLP can be trained with suitable energy predictions, DH cells with a range of domain sizes for each stage should be geometry optimised and their energies compared. This would provide insight as to the stability of a DH cell as a function of domain size. However, no DH cells of any domain size are likely to be lower in energy than the RH cells. DH defects are believed to occur due to kinetic constraints during the transition between stages. Molecular dynamics may provide more insight as to this mechanism but the MLPs used in this work are not capable of modelling charged species, and so a more sophisticated model may be required to accurately reflect electrochemical intercalation. Furthermore, as ML methods and models advance, it may become possible to accurately model phonon spectra. If this occurs, the Raman-active modes for DH cells with larger domain sizes may be compared against experiment as an additional means to determining the DH domain size in experimental cells.

The work in this thesis considers a considerable fraction of hBN defect phase space; expanding far beyond that which has already been modelled in the literature. However, defects were only considered with a neutral charge. This decision was made to constrain the phase space to one that is more practical to explore but charged defects are known to exist and may be responsible for the experimentally-observed SPE. [123, 142, 144] If the optical spectra of the neutral defects generated in this work are not consistent with the experimentally-observed phenomena, the work should be expanded to include charged defects. To avoid conducting further AIRSS searches, the existing database may be used to source defect structures, which may then have a charge applied to the cell. These cells should then be geometry optimised again, and the further optical spectra may be modelled for the low energy charged defects.

The ^{23}Na spectra reported in section 6 show exceptionally good agreement with the experimental spectra collected by Dr Griffith; showing both that ion exchange is an effective method for synthesising alkali-niobates and that the DFT model in use is sufficient for predicting structures and modelling NMR spectra for these systems. It therefore seems appropriate to extend this work to new but related materials. Dr Griffith has already successfully synthesised the Ta analogue of $\text{L-Na}_3\text{NbO}_4$ using the same ion exchange method as applied before. The modelled ^{23}Na spectrum for this material again appears to show good agreement with experiment but data analysis is still ongoing. Furthermore, it has been identified that quenching ion-exchange of Na^+ into $\text{L-Li}_3\text{NbO}_4$ yields a different material to that reported previously. It is believed that this material is the result of incomplete ion-exchange and consequently it may take a form similar to $\text{L-Na}_3\text{NbO}_4$ with some of the Na^+ ions substituted for Li^+ or H^+ . In an approach analogous to structure searching for $\text{NaNb}_7\text{O}_{18}$, a range of $\text{L-Na}_{3-x}\text{H}_x\text{NbO}_4$ and $\text{L-Na}_{3-x}\text{Li}_x\text{NbO}_4$ cells should be generated, geometry optimised, and their ^{23}Na spectra modelled. The spectra may then be compared to experiment to identify the most likely structure. Having verified the identify of these materials, it then becomes easier to assess their physical properties (particularly from a computational approach) such that it may be determined whether they can be implemented in batteries or applied for other applications.

Chapter 9

Supplemental Data

CASTEP [69] Version	19
k-point spacing	0.05 \AA^{-1}
Grid type	Monkhorst-Pack [59, 60]
Pseudopotential	QC5
XC functional	PBE [207]
Cut off energy	650 eV
Mixing scheme	Pulay
Spin polarisation	Off
SEDC scheme	None
Geom op method	LBFGS

Table 9.1 LiNiS geometry optimisation parameters used for initial element swapping and AIRSS. Parameters also used for singlepoints performed on MLP-optimised structures.

Architecture	PaiNN [76]
Cutoff radius	5 \AA
Num. radial basis functions	15
Num. atomic basis sets	128
Num. interactions	3
Batch size	100
Loss function	L1Loss

Table 9.2 LiNiS ANN hyperparameters

CASTEP [69] Version	22
k-point spacing	0.04 \AA^{-1}
Grid type	Monkhorst-Pack [59, 60]
Pseudopotential	USP22
XC functional	PBE [207]
Cut off energy	900 eV
Mixing scheme	Pulay
Spin polarisation	On
SEDC scheme	None
Geom op method	LBFGS

Table 9.3 LiNiS geometry optimisation parameters used for higher accuracy geometry optimisation step

CASTEP [69] Version	21
k-point spacing	0.025 \AA^{-1}
Grid type	Monkhorst-Pack [59, 60]
Pseudopotential	USP19
XC functional	PBE [207]
Cut off energy	900 eV
Mixing scheme	Pulay
Spin polarisation	Off
SEDC scheme	G06 [55]
Geom op method	LBFGS

Table 9.4 KGIC DFT parameters

Architecture	PaiNN [76]
Cutoff radius	5 \AA
Num. radial basis functions	20
Num. atomic basis sets	128
Num. interactions	3
Batch size	50
Loss function	MSELoss

Table 9.5 KGIC ANN hyperparameters

CASTEP [69] Version	19
k-point spacing	0.06 \AA^{-1}
Grid type	Monkhorst-Pack [59, 60]
Pseudopotential	QC5
XC functional	PBESol [208]
Cut off energy	400 eV
Mixing scheme	Pulay
Spin polarisation	On
SEDC scheme	None
Geom op method	LBFGS

Table 9.6 hBN geometry optimisation parameters

CASTEP [69] Version	21
k-point spacing	0.03 \AA^{-1}
Grid type	Monkhorst-Pack [59, 60]
Pseudopotential	USP21
XC functional	PBE [208]
Cut off energy	800 eV
Mixing scheme	Pulay
Spin polarisation	Off
SEDC scheme	None
Geom op method	LBFGS

Table 9.7 NbO geometry optimisation and NMR parameters




Fig. 9.1 Electronic density of states of stage 1-5 RH KGIC and graphite computed with DOS normalised by the number of electrons in the cell and zeroed to the Dirac cone (as determined by the minimum of the DOS). The Fermi level is shown as a horizontal grey bar. C s states are shown in light brown, C p states in dark brown, K s states in light purple, and K p states in dark purple. K-intercalation has a minor impact on the shape of the C states but the Fermi level can be seen shifting to higher energies at greater levels of intercalation. No equivalent figure is provided for DH cells as the Dirac cone is less well-defined. Figure 5.9 shows an alternative representation for RH cells where the DOS is zeroed to the Fermi level. Details of the DFT parameters used can be found in 9.4




Fig. 9.2 Figure shows change in formation energy of defects in response to a change in choice of chemical potentials. X-axis shows a particular set of chemical potentials, with $R\bar{3}$ and $P4_2/nnm$ representing B crystals of those spacegroups, hBN-N being the difference in energy between the chemical potential of hBN and N_2 , $N_2(g)$ being gassesous N_2 , $N_2(s)$ being the ground state molecular crystal of $N_2(s)$, hBN-B being the difference in energy between the chemical potential of hBN and $R\bar{3}$ B, and graphite and diamond being the C allotropes, graphite and diamond. Each connected set of points in the plot shows one of 50 low energy defects. The figure shows that the choice defect formation energy is robust against most choices of chemical potential, with the exception of the hBN-N and hBN-B chemical potentials which treat the formation energy of hBN as 0.




Fig. 9.3 Visualisation of a sample of low energy axial defects and their normalised defect formation energies. Formation energies calculated from chemical potentials of B and C in their ground states. The chemical potential of N is taken as the difference between the energy per formula unit of hBN monolayers and the chemical potential of B. A clearer view of the defects is visible in figure 4.7




Fig. 9.4 Layered Na_3NbO_4 cell with each of the chemical environments shown in figure 6.7 highlighted. Nb shown in dark green, O in red, and Na from environment NaX in purple and other Na ions in yellow.

References

- [1] P. Warde, *Energy Consumption in England & Wales*, Consiglio Nazionale delle Ricerche, 2007, pp. 115–122.
- [2] *Trends in batteries*, <https://www.iea.org/reports/global-ev-outlook-2023/trends-in-batteries>, Accessed: 2023-08-08.
- [3] G. Pistoia, *Advanced Lithium-Ion Batteries for Practical Applications*, Elsevier, 2013, pp. 1–20.
- [4] M. D. Bhatt and C. O'Dwyer, *Phys. Chem. Chem. Phys.*, 2015, **17**, 4799–4844.
- [5] K. Liu, Y. Liu, D. Lin, A. Pei and Y. Cui, *Sci. Adv.*, 2018, **4**, 1–12.
- [6] S. Choi, G. Jung, J. E. Kim, T. Kim and K. S. Suh, *Appl. Surf. Sci.*, 2018, **455**, 367–372.
- [7] F. Wu and G. Yushin, *Energy Environ. Sci.*, 2017, **10**, 435–459.
- [8] S.-H. Yu, X. Feng, N. Zhang, J. Seok and H. D. Abruna, *Acc. Chem. Res.*, 2018, **51**, 273–281.
- [9] J. Ling, S. Hwang, F. Guo, S. Li, Z. Chen, R. Kou, K. Sun, C.-J. Sun, H. Gan, A. Yu, E. A. Stach, H. Zhou and D. Su, *Nat. Commun.*, 2019, **10**, 1–10.
- [10] Y. Ma, *Energy Environ. Mater.*, 2018, **1**, 148–173.

- [11] A. Amardeep, D. J. Freschi, J. Wang and J. Liu, *Nano Res.*, 2022, **1**, 1–28.
- [12] M. Peng, K. Shin, L. Jiang, Y. Jin, K. Zeng, X. Zhou and Y. Tang, *Angew. Chem. Intl. Ed.*, 2022, **61**, 1–22.
- [13] S. H. Lee, M. S. Kim, J.-H. Lee, J.-H. Ryu, V. Do, B. G. Lee, W. Kim and W. I. Cho, *J. Mater. Chem. A*, 2022, **10**, 4157–4169.
- [14] Q. Bai, L. Yang, H. Chen and Y. Mo, *Adv. Energy. Mater.*, 2018, **8**, 1702998.
- [15] J. Xu, Y. Dou, Z. Wei, J. Ma, Y. Deng, Y. Li, H. Liu and S. Dou, *Adv. Sci.*, 2017, **4**, 1700146.
- [16] Y. Li, P. Adelhelm and M.-M. T, *Chem. Soc. Rev.*, 2019, **48**, 4655–4687.
- [17] K. Lu, H. Zhang, F. Ye, W. Luo, H. Ma and Y. Huang, *Energy. Storage. Materials.*, 2019, **16**, 1–5.
- [18] R. Rajagopalan, Y. Tang, X. Ji, C. Jia and H. Wang, *Adv. Funct. Mater.*, 2020, **30**, 1909486.
- [19] J. C. Pramudita, V. K. Peterson, J. A. Kimtpn and N. Sharma, *Powder Diffraction*, 2017, **32**, S43–S48.
- [20] Q. Ran, H. Shi, H. Meng, S.-P. Zeng, W.-B. Wan, W. Zhang, Z. Wen, X.-Y. Lang and Q. Jiang, *Nat. Commun.*, 2022, **13**, 576.
- [21] D. Zagorac, H. Müller, S. Ruehl, J. Zagorac and S. Rehme, *J Appl Crystallogr.*, 2019, **52**, 918–925.
- [22] C. R. Groom, I. J. Bruno, M. P. Lightfofot and S. C. War, *Acta Cryst.*, 2016, **B72**, 171–179.

-
- [23] Y. Zuo, C. Chen, X. Li, Z. Deng, Y. Chen, J. Behler, G. Csányi, A. V. Shapeev, A. P. Thompson, M. A. Wood and S. P. Ong, *J. Phys. Chem. A*, 2020, **124**, 731–745.
- [24] S. Kirkpatrick, C. D. G. Jr. and M. P. Vecchi, *J. Phys. Condens. Matter*, 1983, **220**, 671–680.
- [25] D. V. S. G. Louie, *J. Comp. Phys.*, 1984, **56**, 259–271.
- [26] L. B. Pártay, A. P. Bartók and G. Csányi, *J. Phys. Chem B*, 2010, **114**, 10502–10512.
- [27] C. J. Pickard and R. J. Needs, *J. Phys. Condens. Matter*, 2011, **23**, 053201.
- [28] A. F. Harper, S. P. Emge, P. C. M. M. Magusin, C. P. Grey and A. J. Morris, *Chem. Sci.*, 2023, **14**, 1155–1167.
- [29] D. P. Kingma and J. Ba, *Arxiv*, 2014, **1**, 1412.6980.
- [30] C. G. Broyden, *IMA J Appl Math*, 1970, **6**, 76–90.
- [31] R. Fletcher, *Comput J*, 1970, **13**, 317–322.
- [32] D. Goldfarb, *Math. Comp.*, 1970, **24**, 23–26.
- [33] D. F. Shanno, *Math. Comp.*, 1970, **24**, 647–656.
- [34] E. B. Tadmor, G. S. Smith, N. Bernstein and E. Kaxiras, *Phys. Rev. B*, 1999, **59**, 235–25.
- [35] C. J. Pickard and R. J. Needs, *Phys. Rev. Lett.*, 2006, **97**, 045504.
- [36] A. J. Morris, C. J. Pickard and R. J. Needs, *Phys. Rev. B*, 2008, **78**, 184102.
- [37] A. J. Morris, C. J. Pickard and R. J. Needs, *Phys. Rev. B*, 2009, **80**, 144112.
- [38] J. Mulrone, A. J. Morris and D. M. Duffy, *Phys. Rev. B*, 2011, **84**, 094118.

- [39] A. J. Morris, C. P. Grey, R. J. Needs and C. J. Pickard, *Phys. Rev. B*, 2011, **84**, 224106.
- [40] A. J. Morris, R. J. Needs, E. Salager, C. P. Grey and C. J. Pickard, *Phys. Rev. B*, 2013, **87**, 174108.
- [41] J. Mulroe, M. Watkins, A. J. Morris and D. M. Duffy, *J. Nucl. Mater.*, 2013, **437**, 261–266.
- [42] F. T. Silva, M. X. Silva and J. C. Belchior, *Front. Chem.*, 2019, **7**, 1–21.
- [43] T. C. Le and D. A. Winkler, *Chem. Rev.*, 2016, **116**, 6107–6132.
- [44] M. P. Lourenço, J. Hoastaš, L. B. Herrera, P. Calaminici, A. M. Köster, A. Tchagang and D. R. Salahub, *J Comput Chem.*, 2022, **44**, 814–823.
- [45] J. R. Morris, D. M. Deaven, K. M. Ho, C. Z. Wang, B. C. Pan, J. G. Wacker and D. E. Turner, *Int. At. Energy Agency Bull.*, 1996, **1**, 1–10.
- [46] X. Qu, L. Yang, J. Lv, Y. Xie, J. Yang, Y. Zhang, Y. Wang, J. Zhao, Z. Chen and Y. Ma, *J. Am. Chem. Soc.*, 2022, **144**, 11120–11128.
- [47] S. T. Call, D. Y. Zubarev and A. I. Boldyrev, *J. Am. Chem. Soc.*, 2007, **28**, 1177–1186.
- [48] X. Chen, D. Chen, M. Weng, Y. Jiang, G.-W. Wei and F. Pan, *J. Phys. Chem.*, 2020, **11**, 4392–4401.
- [49] Q. Cui and M. Elstner, *Phys. Chem. Chem. Phys.*, 2014, **16**, 14368–14377.
- [50] E. Schrödinger, *Phys. Rev.*, 1926, **28**, 1049–1070.
- [51] M. Born and R. Oppenheimer, *Ann. Phys.*, 1929, **20**, 457–484.
- [52] J. Kohanoff, *Electronic Structure Calculation for Solids and Molecules*, Cambridge University Press, 2006, pp. 56–57.

- [53] P. Hohenberg and W. Kohn, *Phys. Rev. B*, 1964, **136**, 854–871.
- [54] W. Kohn and L. J. Sham, *Phys. Rev. A*, 1965, **140**, 1133–1138.
- [55] S. Grimme, *J. Comput. Chem*, 2006, **27**, 1788–1799.
- [56] F. Gygi, *J. Chem. Theory Comput.*, 2023, **19**, 1300–1309.
- [57] J. Kohanoff, *Electronic Structure Calculation for Solids and Molecules*, Cambridge University Press, 2006, pp. 143–170.
- [58] J. Kohanoff, *Electronic Structure Calculation for Solids and Molecules*, Cambridge University Press, 2006, pp. 128–143.
- [59] H. J. Monkhorst and J. D. Pack, *Phys. Rev. B*, 1976, **13**, 5188–5192.
- [60] J. D. Pack and H. J. Monkhorst, *Phys. Rev. B*, 1977, **16**, 1748–1749.
- [61] J. Kohanoff, *Electronic Structure Calculation for Solids and Molecules*, Cambridge University Press, 2006, pp. 181–188.
- [62] A. T. Fowler, C. J. Pickard and J. A. Elliott, *J. Phys. Mater.*, 2019, **2**, 034001.
- [63] A. Grisafi, A. Fabrizio, B. Meyer, D. M. Wilkins, C. Corminboeuf and M. Ceriotti, *ACS Cent. Sci.*, 2019, **5**, 57–64.
- [64] A. M. Lewis, A. Grisafi, M. Ceriotti and M. Rossi, *J. Chem. Theory Comput.*, 2021, **17**, 7203–7214.
- [65] H. Yin, Y. Chen and X. Wang, *Matter*, 2022, **5**, 2452–2455.
- [66] J. A. Rckers, L. Tescot, M. Geiger and T. E. Smidt, *Mach. Learn.: Sci. Technol.*, 2023, **4**, 015027.

- [67] P.-Y. Chen, K. Shibata, K. Hagita, T. Miyata and T. Mizoguchi, *J. Phys. Chem. Lett.*, 2023, **14**, 4848–4865.
- [68] P. Pulay, *Chem. Phys. Lett.*, 1980, **73**, 393–398.
- [69] S. J. Clark, M. D. Segall, C. J. Pickard, P. J. Hasnip, M. J. Probert, K. Refson and M. C. Payne, *Zeitschrift fuer Kristallographie*, 2005, **220**, 567–570.
- [70] W. S. McCulloch and W. Pitts, *Bull. math. biophys.*, 1943, **5**, 115–133.
- [71] R. T. Kneusel, *Practical Deep Learning*, No Starch Press, 2021, pp. 200–208.
- [72] J. Behler, *J. Phys. Chem.*, 2011, **134**, 241709.
- [73] M. Gastegger, L. Schwiedrzik, M. Bittermann, F. Berzsenyi and P. Marquetand, *J. Phys. Chem.*, 2018, **148**, 074106.
- [74] K. T. Schütt, P.-J. Kindermans, H. E. Sauceda, S. Chmiela, A. Tkatchenko and K.-R. Müller, *Arxiv*, 2017, **1**, 1706.08566.
- [75] K. T. Schütt, H. E. Sauceda, P.-J. Kindermans, S. Chmiela, A. Tkatchenko and K.-R. Müller, *J. Chem. Phys.*, 2018, **148**, 241722.
- [76] K. T. Schütt, O. T. Unke and M. Gastegger, *Arxiv*, 2021, **1**, 2102.03150.
- [77] K. T. Schütt, P. Kessel, M. Gastegger, K. A. Nicoli, A. Tkatchenko and K.-R. Müller, *J. Chem. Theory Comput.*, 2019, **15**, 448–455.
- [78] K. T. Schütt, S. S. P. Hessmann, N. W. A. Gebauer, J. Lederer and M. Gastegger, *Arxiv*, 2022, **1**, 2212.05517.
- [79] A. H. Larsen, J. J. Mortensen, J. Blomqvist, I. E. Castelli, R. Christensen, M. Dułak, J. Friss, M. N. Groves, B. Hammer, C. Hargus, E. D. Hermes, P. C. Jennings, P. B. Jensen, J. Kermode, J. R. Kitchin, E. L. Kolsbjerg, J. Kubal, K. Kaasbjerg, S. Lysgaard,

- J. B. Maronsson, T. Maxson, T. Olsen, L. Pastewka, A. Peterson, C. Rostgaard, J. Schiøtz, O. Schütt, M. Strange, K. S. Thygesen, T. Vegge, L. Vilhelmsen, M. Walter, Z. Zeng and K. W. Jacobsen, *J. Phys. Condens. Matter*, 2017, **29**, 273002.
- [80] K. Hornik, M. Stinchcombe and H. White, *Neural Netw*, 1989, **2**, 359–366.
- [81] S. H. Simon, *The Oxford Solid State Basics*, Oxford University Press, 2013, pp. 77–95.
- [82] P. Atkins and J. D. Paula, *Atkins' Physical Chemistry*, Oxford University Press, 2006, pp. 513–532.
- [83] A. Manthiram, *Nat. Commun.*, 2020, **11**, 1550.
- [84] J. Molaes, C. Perez-Vicente and J. L. Tirado, *Mat. Res. Bull.*, 1990, **25**, 623–630.
- [85] W. Ebner, D. Fouchard and L. Xie, *Solid State Ion.*, 1994, **69**, 238–256.
- [86] R. Kanno, H. Kubo, Y. Kawamoto, T. Kamiyama, F. Izumi, Y. Takeda and M. Takano, *J. Solid State Chem.*, 1993, **110**, 216–225.
- [87] A. Hirano, R. Kanno, Y. Kawamoto, Y. Takeda, K. Yamaura, M. Takano, K. Ohyama, M. Ohashi and Y. Yamaguchi, *Solid State Ion.*, 1995, **78**, 123–131.
- [88] J. Barker, R. Koksang and M. Y. Saidi, *Solid State Ion.*, 1996, **89**, 25–35.
- [89] R. Stoyanova, E. Zhecheva, R. Alcantara, J. L. Tirado, G. Bromiley, F. Bromiley and T. B. Ballaran, *Solid State Ion.*, 2003, **161**, 197–204.
- [90] H. R. Shaari and V. Sethuprakash, *J. Teknol.*, 2014, **70**, 7–13.
- [91] C. S. Yoon, D.-W. Jun, S.-T. Myung and Y.-K. Sun, *ACS Energy Lett.*, 2017, **2**, 1150–1155.
- [92] Z. Sun, L. Xu, C. Dong, H. Zhang, M. Zhang, Y. Ma, Y. Liu, Z. Li, Y. Zhou, Y. Han and Y. Chen, *Nano Energy*, 2019, **63**, 2211–2855.

- [93] D. Kong, M. Zhang, Y. Xiao, J. Hu, W. Zhao and L. Han, *Elsevier*, 2019, **59**, 327–335.
- [94] L. Biasi, A. Schiele, M. Roca-Ayats, G. Garcia, T. Brezesinski, P. Hartmann and J. Janek, *Chem. Sus. Chem*, 2019, **12**, 2240–2250.
- [95] J. Cen, B. Zhu and D. O. Scanlon, *J. Phys. Energy*, 2023, **5**, 035005.
- [96] J. An, L. Shi, G. Chen, M. Li, H. Li, S. Yuan, S. Chen and D. Zhang, *J. Mater. Chem. A*, 2017, **5**, 19738.
- [97] A. E. Reed and F. Weinhold, *J. Am. Chem. Soc*, 1986, **108**, 3586–3593.
- [98] J. J. Cheng, Y. Ou, J. T. Zhu, H. J. Song and Y. Pan, *Mater. Chem. Phys*, 2019, **231**, 131–137.
- [99] W. Yang, L. Zhou, Z. Luo, J. Zhu, W. Yang, D. Li and L. Fu, *Adv Eng Mater*, 2020, **22**, 1–6.
- [100] H. Ren, J. Wang, Y. Cao, W. Luo and Y. Sun, *Mater. Res. Bull.*, 2021, **133**, 111047.
- [101] N. Feng, D. Hu, P. Wang, X. SUn, X. Li and D. He, *Phys. Chem. Chem. Phys.*, 2013, **15**, 9924–9930.
- [102] S. Ni, X. Yang and T. Li, *Mater. Chem. Phys*, 2012, **132**, 1103–1107.
- [103] H. Ruan, Y. Li, H. Qiu and M. Wei, *J. Alloys Compd.*, 2014, **588**, 357–360.
- [104] J. Wang, S. Y. Chew, D. Wexler, G. X. Wang, S. H. Ng, Z. Zhong and H. K. Liu, *Electrochem. commun.*, 2007, **9**, 1877–1880.
- [105] A. J. Morris, C. P. Grey and C. J. Pickard, *Phys. Rev. B*, 2014, **90**, 054111.
- [106] X. Xu, Y. Z. H. You, H. Min, X. Wu, W. Zang, J. Hao, Z. Feng, X. Liu and H. Yang, *Appl. Surf. Sci.*, 2022, **605**, 154758.

- [107] X. J. Liu, Z. Z. Xu, H. J. Ahn, S. K. Ahn, S. K. Lyu and I. S. Ahn, *Powder Technol.*, 2012, **229**, 24–29.
- [108] C. Jin, L. Zhou, L. Fu, J. Zhu, D. Li and W. Yang, *J. Power Sources*, 2017, **352**, 83–89.
- [109] J. Wi, S. Liu, Y. Rehman, T. Huang, J. Zhao, Q. Gu, J. Mao and Z. Guo, *Adv. Funct. Mater.*, 2021, **31**, 2010832.
- [110] W. Li, B. Huang, Z. Liu, J. Yang, Y. Li, S. Xiao, Q. Chen, G. Li, X. Zhao and W. Zhang, *Electrochimica Acta.*, 2021, **369**, 4686.
- [111] K. Li, J. Yu, Q. Wang, L. Li, W. Zhang, J. Ma, J. Zhang, P. Liu and D. Li, *Nanotechnology*, 2021, **32**, 195406.
- [112] K. Han, J. Meng, X. Hong, X. Wang and L. Mai, *Nanoscale*, 2020, **12**, 8255.
- [113] J. Zhang, Y. Li, M. Han, Q. Xia, Q. Chen and M. Chen, *Mater. Res. Bull.*, 2021, **137**, 111186.
- [114] S. Wang, P. Ning, S. Huang, W. Wang, S. Fei, Q. He, J. Zai, Y. Jiang, Z. Hu, X. Qian and Z. Chen, *J. Power Sources*, 2019, **436**, 226857.
- [115] X. Li, M. Banis, A. Lushinton, X. Yang, Q. Sun, Y. Zhao, C. Liu, Q. Li, B. Wang, W. Xiao, C. Wang, M. Li, J. Liang, R. Li, Y. Hu, L. GOncharova, H. Zhang, T.-K. Sham and X. Sun, *Nat. Commun.*, 2018, **9**, 4509.
- [116] C. Luo, E. Hu, K. J. Gaskell, X. Fan, T. Gao, C. Cui, S. Ghose, X.-Q. Yang and C. Wang, *Proc. Natl. Acad. Sci. U.S.A.*, 2018, **117**, 14712–14720.
- [117] A. Jain, S. P. Ong, G. Hautier, W. Chen, W. D. Richards, S. Dacek, S. Cholia, D. Gunter, D. Skinner, G. Ceder and K. A. Persson, *APL Materials*, 2013, **220**, 011022.
- [118] S. Kirklin, J. E. Saal, B. Meredig, A. Thompson, J. W. Doak, M. Aykol, S. Rühl and C. Wolverton, *NPJ Comput. Mater.*, 2015, **1**, 15010.

- [119] A. F. Harper, M. L. Evans and A. J. Morris, *Chem. Mater.*, 2020, **32**, 6629–6639.
- [120] L. E. marbella, M. L. Evans, M. F. Groh, J. Nelson, K. J. Griffith, A. J. Morris and C. P. Grey, *J. Am. Chem. Soc.*, 2018, **140**, 7994–8004.
- [121] M. Mayo and A. J. Morris, *Chem. Mater.*, 2017, **29**, 5787–5795.
- [122] C. George, A. J. Morris, M. H. Hodarres and M. D. Volder, *Chem. Mater.*, 2016, **28**, 7304–7310.
- [123] A. B. D. a.-j.-w.-i. Shaik and P. Palla, *Sci.*, 2021, **11**, 12285.
- [124] T. T. Tran, K. Bray, M. J. Ford, M. Tothh and I. Aharonovich, *Nat. Nanotechnol.*, 2015, **11**, 37–42.
- [125] M. Hijlkema, B. Weber, H. P. Specht, S. C. Webster, A. Kuhn and G. Rempe, *Nat. Phys.*, 2007, **3**, 253–255.
- [126] M. Steiner, A. Hartschuh, R. Korlacki and A. J. Meixner, *Appl. Phys. Lett.*, 2007, **90**, 183122.
- [127] V. Zwiller, T. Aichele, W. Seifert, J. Persson and O. Benson, *Appl. Phys. Lett.*, 2003, **82**, 1509–1511.
- [128] X. Ma, N. F. Hartmann, J. K. S. Baldwin, S. K. Doorn and H. Htoon, *Nat. Nanotechnol.*, 2015, **10**, 671675.
- [129] J. C. Bienfang, J. Fan, A. Migdall and S. V. Polyakov, *Single-photon generation and detection : physics and applications*, Elsevier, 2013, pp. 1–24.
- [130] V. Giovannetti, S. Lloyd and K. Maccone, *Nat. Photonics*, 2011, **5**, 222–229.
- [131] A. Aspuru-Guzik and P. Walther, *Nat. Phys.*, 2012, **8**, 285–291.

- [132] S. Castelletto, B. C. Johnson, V. Ivády, N. Stavrias, T. Umeda, A. Gali and T. Ohshima, *Nat. Mater.*, 2014, **13**, 151–156.
- [133] H.-K. Lo, M. Curty and K. Tamaki, *Nat. Photonics*, 2014, **8**, 595–604.
- [134] T. Zhong, J. M. Kindem, E. Miyazono and A. Faraon, *Nat. Photonics*, 2015, **6**, 1–6.
- [135] D. Rosenberg, A. J. Kerman, R. J. Molnar and E. A. Dauler, *Nat. Photonics*, 2013, **21**, 1440–1447.
- [136] M. Bozzio, M. Vyvlecka, M. Cosacchi, C. Nawrath, T. Sidelmann, J. C. Loredó, S. L. Portalupi, V. M. Axt, P. Michler and P. Walther, *npj Quantum Inf*, 2022, **8**, 104.
- [137] S. A. Tawfik, S. Ali, M. Fronzi, M. Kianinia, T. T. Tran, C. Stampfl, I. Aharonovich, M. Toth and M. J. Ford, *Nanoscale*, 2017, **9**, 13575–13582.
- [138] N. P. Blascetta, M. Liebel, X. Lu, T. Taniguchi, K. Watanabe, D. K. Efetov and N. F. van Hulst, *Nano Lett.*, 2020, **20**, 1992–1999.
- [139] J. Strand, L. Larcher and A. L. Shluger, *J. Phys.*, 2020, **32**, 1–10.
- [140] Y. A. Korkmaz, C. Bulutay and C. Sevik, *Physica B Condens. Matter*, 2020, **584**, 4119599.
- [141] M. E. Turiansky, A. Alkauskas and C. G. V. de Walle, *Nat. Mater.*, 2020, **19**, 481–490.
- [142] N. Mendelson, D. Chugh, J. R. Reimers, T. S. Cheng, A. Gottscholl, H. Long, C. J. Mellor, A. Zettl, V. Dyakonov, P. H. Beton, S. V. Novikov, C. Jagadish, H. H. Tan, M. J. Ford, M. Toth, C. Bradac and I. Aharonovich, *Nat. Mater.*, 2021, **20**, 321–328.
- [143] D. S. Wang, C. J. Ciccarino, J. Flick and P. Narang, *ACS Nano*, 2021, **15**, 5240–5248.
- [144] P. Auburger and A. Gali, *Phys. Rev. B*, 2021, **104**, 075410.

- [145] M. Andrea, H. Conor, G. Myrta and V. Daniele, *Comput Phys Commun*, 2009, **180**, 1392–1403.
- [146] F. Aryasetiawan and O. Gunnarsson, *Rep. Prog. Phys.*, 1998, **61**, 237.
- [147] E. E. Salpeter and H. A. Bethe, *Phys. Rev.*, 1951, **84**, 1232–1241.
- [148] J. R. Reimers, A. Sajid, R. Kobayashi and M. J. Ford, *J. Chem. Theory Comput.*, 2018, **14**, 1602–1613.
- [149] L. Weston, D. Wickramaratne, M. Mackoiti, A. Alkauskas and C. G. V. de Walle, *Phys. Rev. B*, 2018, **97**, 214104.
- [150] K. Feidenhagen and G. Cadenbach, *Z. Anorg. Allg. Chem.*, 1926, **158**, 249–263.
- [151] V. A. Schleede and H. Wellmann, *Z. physikal. Chem. Abt. B. Bd. 18, Heft*, 1929, **178**, 1–28.
- [152] K. Mizushima, P. C. Jones, P. J. Wiseman and J. Goodenough, *Mat. Res. Bull.*, 1980, **15**, 783–789.
- [153] R. Yazami and P. H. Touzain, *J. Power Sources*, 1982, **9**, 365–371.
- [154] W. M. Haynes, *CRC Handbook of Chemistry and Physics*, CRC Press, 97th edn., 2016, pp. 14–17.
- [155] J. Xu, Y. Dou, Z. Wei, J. Ma, Y. Deng, Y. Li, H. Liu and S. Dou, *Adv. Sci.*, 2017, **4**, 1700146.
- [156] Y. Li, Y. Lu, P. Adelhelm, M. M. T. and Y. S. Hu, *Chem. Soc. Rev.*, 2019, **48**, 4655–4687.
- [157] Q. Bai, L. Yang, H. Chen and Y. Mo, *Adv. Energy Mater.*, 2018, **8**, 1702998.

- [158] R. Rajagopalan, Y. Tang, X. Ji, C. Jia and H. Wang, *Adv. Funct. Mater.*, 2020, **30**, 1909486.
- [159] K. M. Abraham, *ACS Energy Lett.*, 2020, **5**, 3544–3547.
- [160] C. Zhao, Y. Lu, L. Chen and Y.-S. Hu, *Info. Mat*, 2019, **2**, 126–138.
- [161] K. Nobuhara, H. Nakayama, M. Nose, Shinji, Nakanishi and H. Iba, *J. Power Sources*, 2013, **243**, 585–587.
- [162] W. Rüdorff, *Adv. Inorg. Chem*, 1959, **1**, 223–266.
- [163] J. C. Chacón-Torres, L. Wirtz and T. Pichler, *Physica Status Solidi (B) Basic Research*, 2014, **251**, 2337–2355.
- [164] W. Luo, J. Wan, B. Ozdemir, W. Bao, Y. Chen, J. Dai, H. Lin, Y. Xu, F. Gu, V. Barone and L. Hu, *Nano Lett.*, 2015, **15**, 7671–7677.
- [165] J. Zhao, X. Zou, Y. Zhu, Y. Xu and C. Wang, *Adv. Funct. Mater.*, 2016, **26**, 8103–8110.
- [166] E. Zhou, J. Xi, Y. Liu, Z. Xu, Y. Guo, L. Peng, W. Gao, J. Ying, Z. Chen and C. Gao, *Nanoscale.*, 2017, **9**, 18613–18618.
- [167] J. Liu, T. Yin, B. Tian, B. Zhang, C. Qian, Z. Wang, L. Zhang, P. Liang, Z. Chen, J. Yan, X. Fan, J. Lin, X. Chen, Y. Huang, K. P. Loh and Z. X. Shen, *Adv. Energy Mater.*, 2019, **9**, 1900579.
- [168] V. W. Rüdorff and U. Hofmann, *Z. Anorg. Allg. Chem.*, 1938, **238**, 1–50.
- [169] C. D. Fuerst, J. E. Fischer, J. D. Axe, J. B. Hastings and D. B. M. Whan, *Phys. Rev. Lett*, 1983, **50**, 357–360.
- [170] Y. Wang, P. Puech, I. Gerber and A. Pénicaud, *J Raman Spectrosc*, 2014, **45**, 219–223.

- [171] H. Onuma, K. Kubota, S. Muratsubaki, W. Ota, M. Shishkin, H. Sato, K. Yamashita, S. Yasuno and S. Komaba, *J. Mater. Chem. A*, 2020, **9**, 11187–11200.
- [172] M. D. M. Rahman, C. Hou, S. Mateti, K. Tanwar, I. Sultana, A. M. Glushenkov and Y. Chen, *Power Sources*, 2020, **476**, 228733.
- [173] J. Zhao, X. Zou, Y. Zhu, Y. Xu and C. Wang, *Adv. Funct. Mater.*, 2016, **26**, 8103–8110.
- [174] J. C. Chacón-Torres, L. Wirtz and T. Pichler, *ACS Nano*, 2013, **7**, 9249–9259.
- [175] C. A. Howard, M. P. M. Dean and F. Withers, *Phys. Rev. B.*, 2011, **84**, 241404.
- [176] U. Fano, *Phys. Rev.*, 1961, **124**, 1866–1878.
- [177] A. C. Ferrari, *Solid State Commun.*, 2007, **143**, 47–57.
- [178] A. C. Ferrari and D. M. Basko, *Nat. Nanotechnol.*, 2013, **8**, 235–246.
- [179] K. Share, A. P. Cohn, R. E. Carter and C. L. Pint, *Nanoscale*, 2016, **8**, 16435–16439.
- [180] A. C. Ferrari, *Phys. Rev. Lett.*, 2006, **97**, 187401.
- [181] R. J. Nemanich, S. A. Solin and D. Guerard, *Phys. Rev. B.*, 1977, **16**, 2965–2973.
- [182] J. E. Fischer and T. E. Thompson, *Physics Today*, 1978, **31**, 36–44.
- [183] S. Wang, C. Wang and X. Ji, *RSC Adv.*, 2017, **7**, 52252–52260.
- [184] M. M. Rahman, C. Hou, S. Mateti, K. Tanwar, I. Sultana, A. M. Glushenkov and . Chen, *J. Power Sources*, 2020, **476**, 228733.
- [185] P. Vecera, J. C. Chacón-Torres, T. Pichler, S. Reich, H. R. Soni, A. Görling, K. Edelthammer, H. Peterlik, F. Hauke and A. Hirsch, *Nat. Commun.*, 2017, **8**, 1–9.
- [186] D. Tristant, Y. Wang, I. C. Gerber, M. Monthieux, A. Pénicaud, P. Puech, I. Gerber and A. P. Enicaud, *J. Appl. Phys.*, 2015, **118**, 44304.

- [187] N. Daumas and M. A. Hérold, *C. R. Acad. Sc. Paris*, 1969, **268**, 373–375.
- [188] M. S. Dresselhaus and G. Dresselhaus, *Intercalation compounds of graphite*, Adv. Physics, 2002, pp. 31–35.
- [189] P. G. de Gennes and J. Prost, *The Physics of Liquid Crystals*, Oxford University Press, 1995, p. 616.
- [190] J. J. Murray and A. R. Ubbelohde, *Proc. Roy. Soc. A.*, 1969, **312**, 371–380.
- [191] J. Zhang, Y. Cheng, W. Lu, E. Briggs, A. J. Ramirez-Cuesta and J. Bernholc, *J. Chem. Theory Comput.*, 2019, **15**, 6859–6864.
- [192] B. G. Blasse, *Z Anorg Allg Chem*, 1963, **326**, 44–46.
- [193] K. Ukei, H. Suzuki, T. Shishido and T. Fukuda, *Acta. Cryst*, 1994, **50**, 655–656.
- [194] N. Yabuuchi, M. Takeuchi, M. Nakayama, H. Shiiba, M. Ogawa, K. Nakayama, T. Ohta, D. Endo, T. Ozaki, T. Inamasu, K. Sato and S. Komaba, *Proc. Natl. Acad. Sci. U.S.A.*, 2015, **112**, 7650–7655.
- [195] V. L. McLaren, C. A. Kirk, M. Poisot, M. Castellanos and A. R. West, *Dalton Trans.*, 2004, **7**, 3042–3047.
- [196] N. Yabuuchi, *Chem. Rec.*, 2019, **19**, 690–707.
- [197] X. Fan, Q. Qin, D. Liu, A. Dou, M. Su, Y. Liu and J. Pan, *J. Alloys Compd*, 2019, **797**, 961–969.
- [198] N. Kuganathan, A. Kordatos, N. Kelaidis and A. Chroneos, *Sci. Rep.*, 2019, **9**, 2192.
- [199] G. Blasse and G. P. M. van de Heuval, *Mat. Res. Bull.*, 1972, **7**, 1041–1044.
- [200] A. P. Pivovarova, *Glass Phys. Chem.*, 2004, **30**, 274–279.

- [201] S. Yamozoe, K. Zhibata, K. Kato and T. Wada, *Chem. Lett.*, 2013, **42**, 380–382.
- [202] K. Sato, M. Nakayama, A. M. Glushenkov, T. Mukai, Y. Hashimoto, K. Yamanaka, M. Yoshimura, T. Ohta and N. Yabuuchi, *Chem. Mater.*, 2017, **29**, 5043–5047.
- [203] H. Ren, Y. Li, Q. Ni, Y. Bai, H. Zhao and C. Wu, *adv. Mater.*, 2022, **34**, 2106171.
- [204] H. Han, Q. Jacquet, Z. Jiang, F. N. Sayed, J.-C. Jeon, A. Sharma, A. M. Shankler, A. Kakekhani, H. L. Meyerheim, J. Park, S. Y. Nam, K. J. Griffith, L. Simonelli, A. M. Rappe, C. P. Grey and S. S. P. Parkin, *Nat. Mater.*, 2023, **22**, 1128–1135.
- [205] B.-O. Marinder and M. Sundberg, *Acta Cryst.*, 1984, **40**, 82–86.
- [206] T. G. Bergmann, M. O. Welzel and C. R. Jacob, *Chem. Sci*, 2020, 1862–1877.
- [207] J. P. Perdew, K. Burke and M. Ernzerhof, *Phys. Rev. Lett.*, 1996, **77**, 3865–3868.
- [208] J. P. Perdew, A. Ruzsinszky, G. I. Csonka, O. A. Vydrov, G. E. Scuseria, L. A. Constantin, X. Zhou and K. Burke, *Phys. Rev. Lett.*, 2008, **100**, 136406.